



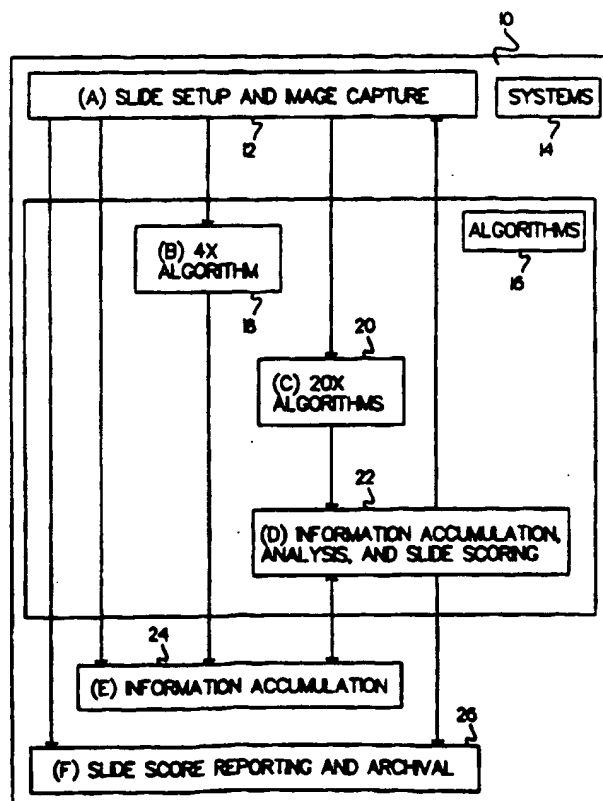
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06K 9/00	A1	(11) International Publication Number: WO 96/09598 (43) International Publication Date: 28 March 1996 (28.03.96)
<p>(21) International Application Number: PCT/US95/10165</p> <p>(22) International Filing Date: 11 September 1995 (11.09.95)</p> <p>(30) Priority Data: 08/309,931 20 September 1994 (20.09.94) US</p> <p>(71) Applicant: NEOPATH, INC. [US/US]; 8271 154th Avenue, N.E., Redmond, WA 98052 (US).</p> <p>(72) Inventors: LEE, Shih-Jong, J.; 14116 Southeast 46th Street, Bellevue, WA 98006 (US). WILHELM, Paul, S.; 8217 122nd Avenue, N.E., Kirkland, WA 98033 (US). MEYER, Michael, G.; 201 West Olympic Place, No. 301, Seattle, WA 98119 (US). BANNISTER, Wendy, R.; 358 North 82nd Street, Seattle, WA 98103 (US). KUAN, Chih-Chau; 18217 N.E. 100th Court, Redmond, WA 98052 (US). ORTYN, William, E.; 27531 N.E. 149th Lane, Devall, WA 98019 (US). NELSON, Larry, A.; 15930 N.E. 6th Street, Bellevue, WA 98008 (US). FROST, Keith, L.; 15310 26th Avenue, N.E., Seattle, WA 98115 (US). HAYENGA, Jon, W.; 21648 94th Place South, Kent, WA 98031 (US).</p> <p>(74) Agent: LEONE, George, A., Sr.; Leone & Moffa, P.A., Suite 1275, 601 Lakeshore Parkway, Minneapolis, MN 55305 (US).</p>		<p>(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: **CYTOLOGICAL SLIDE SCORING APPARATUS**

(57) Abstract

The invention detects areas of interest at low magnification (18), locating possible abnormal cells or other cells of interest using image processing and statistical pattern recognition techniques. Next, at high magnification (20), the areas identified at low magnification are re-examined. The information from low magnification (18) and high magnification (20) scans is collated (24) and a determination is made about the slide - whether it is normal, abnormal, contains endocervical component, and so forth. The invention also provides a method and apparatus to train object feature and slide feature classifiers. The invention provides an automated cytology system that can process training slides for use in the feed back classifier development environment. The invention also can classify endocervical group of cells.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

- 1 -

CYTOLOGICAL SLIDE SCORING APPARATUS

The invention relates to an automated cytology system and more particularly to an automated Papanicolaou-stained gynecological smear analysis system.

BACKGROUND OF THE INVENTION

The currently well established human review process for Papanicolaou smear analysis follows standards recommended by the Bethesda System. A trained cytotechnologist systematically views a slide at low magnification to identify areas of interest. When an area of interest is identified, the cytotechnologist views the area at high magnification where it is possible to distinguish abnormal cells according to changes in their structure and context.

In much the same way as a human reviews Papanicolaou smears, it is one motivation of the invention to scan slides at low magnification to detect possible areas of interest, then at high magnification, to scan those areas to locate possible abnormal cells or other cells of interest. As a cytotechnologist compares size, shape and density of cells against established criteria, it is an additional motive of the invention to compare objects according to criteria established during a training process.

SUMMARY OF THE INVENTION

The invention detects areas of interest at low magnification, locating possible abnormal cells or other cells of interest using image processing and statistical pattern recognition techniques. Next, at high magnification, the areas identified at low magnification are re-examined. The information from the low magnification and high magnification scans is collated and a determination is made about the slide

- 2 -

- whether it is normal, abnormal, contains endocervical component, and so forth. The invention also provides a method and apparatus to train object feature and slide feature classifiers. The invention
5 provides an automated cytology system that can process training slides for use in a feedback classifier development environment. The invention also can classify endocervical groups of cells.

Other objects, features and advantages of the
10 present invention will become apparent to those skilled in the art through the description of the preferred embodiment, claims and drawings herein wherein like numerals refer to like elements.

BRIEF DESCRIPTION OF THE DRAWINGS

15 To illustrate this invention, a preferred embodiment will be described herein with reference to the accompanying drawings.

Figures 1A, 1B and 1C show the apparatus of the invention.

20 Figure 2A shows the classification method of the invention.

Figure 2B shows an information flow diagram of the method of the invention.

25 Figure 3 shows the slide processing method of the invention.

Figure 4 shows the 4x magnification scan processing of the invention.

Figure 5 shows the 20x magnification scan processing of the invention.

30 Figure 6 shows the apparatus of the invention with system software interfacing.

Figure 7 shows a schematic diagram of the classifier development method of the invention.

35 Figure 8 shows one example of a classifier architecture of the method of the invention.

- 3 -

Figure 9 shows a classifier training data flow diagram of the method of the invention.

Figure 10 shows an object classifier training process flow diagram of the method of the invention.

5 Figure 11 shows a semi-automatic image processing method.

Figure 12 shows the classifier construction method of the invention.

10 Figure 13 shows the classifier development process flow diagram of the method of the invention.

Figure 14 shows the slide evaluation process flow diagram of the method of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

15 In a presently preferred embodiment of the invention, the system disclosed herein is used in a system for analyzing cervical pap smears, such as that shown and disclosed in U.S. Patent Application Serial No. 07/838,064, entitled "METHOD FOR IDENTIFYING NORMAL BIOMEDICAL SPECIMENS", by Alan C. Nelson, et al., filed February 18, 1992; U.S. Patent Application
20 Serial No. 08/179,812 filed January 10, 1994 which is a continuation in part of U.S. Patent Application Serial No. 07/838,395, entitled "METHOD FOR IDENTIFYING OBJECTS USING DATA PROCESSING TECHNIQUES",
25 by S. James Lee, et al., filed February 18, 1992; U.S. Patent Application Serial No. 07/838,070, now U.S. Pat. No. 5,315,700, entitled "METHOD AND APPARATUS FOR RAPIDLY PROCESSING DATA SEQUENCES", by Richard S. Johnston, et al., filed February 18, 1992; U.S. Patent
30 Application Serial No. 07/838,065, filed 02/18/92, entitled "METHOD AND APPARATUS FOR DYNAMIC CORRECTION OF MICROSCOPIC IMAGE SIGNALS" by Jon W. Hayenga, et al.; and U.S. Patent Application Serial No. 08/302,355, filed September 7, 1994 entitled "METHOD
35 AND APPARATUS FOR RAPID CAPTURE OF FOCUSED MICROSCOPIC

- 4 -

IMAGES" to Hayenga, et al., which is a continuation-in-part of Application Serial No. 07/838,063 filed on February 18, 1992 the disclosures of which are incorporated herein, in their entirety, by the foregoing references thereto.

The present invention is also related to biological and cytological systems as described in the following patent applications which are assigned to the same assignee as the present invention, filed on September 20, 1994 unless otherwise noted, and which are all hereby incorporated by reference including U.S. Patent Application No. 08/309,405 to Frost, et al., entitled "AUTOMATIC FOCUSING OF BIOMEDICAL SPECIMENS APPARATUS," U.S. Patent Application Serial No. 08/309,064 to Ortyn, et al., entitled "APPARATUS FOR ILLUMINATION STABILIZATION AND HOMOGENIZATION," U.S. Patent Application No. 08/309,249 to Ortyn et al., entitled "BIOLOGICAL SPECIMEN ANALYSIS SYSTEM PROCESSING INTEGRITY," U.S. Patent Application Serial No. 08/303,179, filed September 8, 1994, to Ortyn et al., entitled "CYTOLOGICAL SYSTEM ILLUMINATION INTEGRITY CHECKING APPARATUS AND METHOD," U.S. Patent Application Serial No. 08/309,078, to Ortyn et al., entitled "CYTOLOGICAL SYSTEM IMAGE COLLECTION INTEGRITY CHECKING APPARATUS," U.S. Patent Application Serial No. 08/309,130 to Ortyn et al., entitled "CYTOLOGICAL SYSTEM AUTOFOCUS INTEGRITY CHECKING APPARATUS," U.S. Patent Application Serial No. 08/308,140, filed September 19, 1994, to Ortyn et al., entitled "AUTOMATED CYTOLOGICAL SYSTEM POSITION INTEGRITY CHECKING METHOD AND APPARATUS," U.S. Patent Application Serial No. 08/309,118, to Kuan et al., entitled, "FIELD PRIORITIZATION APPARATUS AND METHOD," U.S. Patent Application Serial No. 08/309,061, to Wilhelm et al., entitled "APPARATUS FOR AUTOMATED

- 5 -

IDENTIFICATION OF CELL GROUPINGS ON A BIOLOGICAL SPECIMEN," U.S. Patent Application Serial No. 08/309,116 to Meyer et al. entitled "APPARATUS FOR AUTOMATED IDENTIFICATION OF THICK CELL GROUPINGS ON A BIOLOGICAL SPECIMEN," U.S. Patent Application Serial No. 08/309,115 to Lee et al. entitled "BIOLOGICAL ANALYSIS SYSTEM SELF CALIBRATION APPARATUS," U.S. Patent Application Serial No. 08/308,992, to Lee et al. entitled "APPARATUS FOR IDENTIFICATION AND INTEGRATION OF MULTIPLE CELL PATTERNS," U.S. Patent Application Serial No. 08/309,063 to Lee et al. entitled "A METHOD FOR CYTOLOGICAL SYSTEM DYNAMIC NORMALIZATION," U.S. Patent Application Serial No. 08/309,248 to Rosenlof et al. entitled "METHOD AND APPARATUS FOR DETECTING A MICROSCOPE SLIDE COVERSIP," U.S. Patent Application Serial No. 08/309,077 to Rosenlof et al. entitled "APPARATUS FOR DETECTING BUBBLES IN COVERSIP ADHESIVE," U.S. Patent Application Serial No. 08/309,148 to Lee et al. entitled "METHOD AND APPARATUS FOR IMAGE PLANE MODULATION PATTERN RECOGNITION," U.S. Patent Application Serial No. 08/309,250 to Lee et al. entitled "APPARATUS FOR THE IDENTIFICATION OF FREE-LYING CELLS," U.S. Patent Application Serial No. 08/309,209 to Oh et al. entitled "A METHOD AND APPARATUS FOR ROBUST BIOLOGICAL SPECIMEN CLASSIFICATION," U.S. Patent Application Serial No. 08/309,117, to Wilhelm et al. entitled "METHOD AND APPARATUS FOR DETECTION OF UNSUITABLE CONDITIONS FOR AUTOMATED CYTOLOGY SCORING."

It is to be understood that the various processes described herein may be implemented in software suitable for running on a digital processor. The software may be embedded, for example, in the central processor 540.

- 6 -

The invention utilizes a set of image analysis and decision making apparatuses. Their purpose is to provide information for the control of slide processing and to provide a cytologic evaluation for each slide analyzed.

Now refer to Figures 1A, 1B and 1C which show a schematic diagram of one embodiment of the apparatus of the invention. The apparatus of the invention comprises an imaging system 502, a motion control system 504, an image processing system 536, a central processing system 540, and a workstation 542. The imaging system 502 is comprised of an illuminator 508, imaging optics 510, a CCD camera 512, an illumination sensor 514 and an image capture and focus system 516. The image capture and focus system 516 provides video timing data to the CCD cameras 512, the CCD cameras 512 provide images comprising scan lines to the image capture and focus system 516. An illumination sensor intensity is provided to the image capture and focus system 516 where an illumination sensor 514 receives the sample of the image from the optics 510. In one embodiment of the invention, the optics may further comprise an automated microscope. The illuminator 508 provides illumination of a slide. The image capture and focus system 516 provides data to a VME bus 538. The VME bus distributes the data to an image processing system 536. The image processing system 536 is comprised of field-of-view processors 568. The images are sent along the image bus 564 from the image capture and focus system 516. A central processor 540 controls the operation of the invention through the VME bus 538. In one embodiment the central processor 562 comprises a Motorola 68030 CPU. The motion controller 504 is comprised of a tray handler 518, a microscope stage controller 520, a microscope tray

- 7 -

controller 522, and a calibration slide 524. The motor drivers 526 position the slide under the optics. A bar code reader 528 reads a barcode located on the slide 524. A touch sensor 530 determines whether a slide is under the microscope objectives, and a door interlock 532 prevents operation in case the doors are open. Motion controller 534 controls the motor drivers 526 in response to the central processor 540. An Ethernet communication system 560 communicates to a workstation 542 to provide control of the system. A hard disk 544 is controlled by workstation 550. In one embodiment, workstation 550 may comprise a Sun Spark Classic (TM) workstation. A tape drive 546 is connected to the workstation 550 as well as a modem 548, a monitor 552, a keyboard 554, and a mouse pointing device 556. A printer 558 is connected to the Ethernet 560.

The central computer 540 runs a real time operating system, controls the microscope and the processor to acquire and digitize images from the microscope. The flatness of the slide may be checked, for example, by contacting the four corners of the slide using a computer controlled touch sensor. The computer 540 also controls the microscope stage to position the specimen under the microscope objective, and from one to 15 field of view (FOV) processors 568 which receive images under control of the computer 540.

Initially the computer 540 implements a magnification 4x analysis of the slide providing information for the control of further slide processing. The 4x analysis extracts cellular information from low magnification (4x) images. The 4x image is divided into a number of 20x images. Those skilled in the art will recognize that the low

- 8 -

resolution and high resolution magnifications may differ without deviating from the scope of the invention. For each 20x region of a 4x image, scores are generated representing the likelihood that the 20x region (also known as 20x field of view or 20x FOV) contains abnormal cells or cellular groupings. After scores have been generated for each 20x FOV on the slide during the 4x scan, a 20x FOV selection analysis uses the scores to prioritize the 20x scan.

10 It is to be understood that the various processes described herein may be implemented in software suitable for running on a digital processor. The software may be embedded, for example, in the central processor 540.

15 At high magnification, several different analyses are used to extract information from images; several are used because abnormality may be present in different forms. Additionally, cellular abnormality is not the only condition that the invention identifies as needing review. If the cell sample was collected inadequately, or preserved or stained improperly, the slide may contain obscured, degenerate, or incomplete cellular material. In one embodiment of the invention, laboratory personnel review such slides as well as those where abnormality has been identified. The following analyses are performed to extract information from the 20x magnification images: single cell analysis, group analysis, thick group analysis, and stripe analysis.

25 Slide scores are calculated from data gathered during the 4x and 20x scans. The computer 540 generates scores after slide scoring analysis such as quality control score (QC score), endocervical score, slide suitability score, and cell count score.

35 The computer 540 performs two basic operations:

- 9 -

information extraction and information analysis. Examples of information extraction operations are the 4x magnification analysis and 20x magnification analysis, which analyze images. Examples of
5 information analysis operations are the slide scoring analysis and the 20x FOV selection analysis, which make decisions based on extracted data.

The information extraction operation examines images and calculates measures of image content.
10 Generally, information about objects in each image are of the greatest interest.

The process of detecting objects and calculating measures or features of the objects is shown in Figure 2A. The object detection process 6 detects objects
15 through image segmentation. These objects represent potential objects of interest; for example, an abnormal cell nucleus.

The feature calculation process 7 measures properties such as perimeter or area for each detected
20 object. After feature calculation 7, the invention compares 8 the feature value for an unknown object to the feature values for known objects 4, allowing the unknown object to be classified 9. The particular feature representation and classification methods
25 determine the comparison criterion.

The information analysis process takes information gathered from various sources and produces specific results. For example, the slide scoring
30 process takes information from the information extraction operation and calculates scores that represent cytologic evaluations of a slide.

Computer 540 controls slide analysis by commanding a number of specialized supporting
processors 568. Image information extraction is
35 handled by FOV processors. Each FOV processor 568 has

- 10 -

an image queue, so computer 540 can sequence images to each FOV processor while information extraction continues on images previously placed in the queue. Once an FOV processor 568 has completed the analysis of an image, it passes information about the image to computer 540 where the information is stored for later analysis and slide scoring.

Because each FOV processor 568 has access only to information from one image at a time, and because computer 540 has access to information from the entire slide, the host processor must complete classification for detected objects that require multi-image features. As a result, all object detection is done on the FOV processors 568, but not all feature calculation and object classification is done on the FOV processors 568. Some intermediate results must be passed from the FOV processor 568 to computer 540 for later completion of feature calculation and object classification. Figure 2A shows the generic flow diagram for image information extraction.

Feature calculation 7 and object classification 8 may happen on the FOV processor 568 or on both the FOV processor 568 and the host processor 540. If multi-image features are not required for object classification 8, feature calculation 7 and object classification will not occur on the host processor 540. If multi-image features are required, then feature calculation and object classification will occur on both the FOV processor 568 and the computer 540. Those skilled in the art will recognize that when multi-scan images are analyzed processing of all images cannot be completed until the information needed to calculate multi-image features is available.

Figure 2B shows the invention's slide handling

- 11 -

software organized as system software 14 and analysis software 16. The diagram represents logical associations between subsystems within analysis software and system software. Note that system software has been broken into three subsystems:

- 5 ◦ Slide Setup and Image Capture 12, primarily provides images to the analysis software.
- Information Accumulation 22, accumulates slide data during processing.
- 10 ◦ Slide Score Reporting and Archival 26, organizes and reports slide results to the user.

Arrows in Figure 2B represent information flow, including image data, object features, and analysis results. The invention provides a slide suitability score for biological specimen slides. The slide suitability score is one of six slide scores that are the results of processing slides on the automated cytology quality control system of the invention.

15 The slide suitability score is derived from measurements of slide characteristics and machine effectiveness. Slide characteristics are properties of the sample and slide, such as bubble area, cover slip and mounting medium thickness, and staining properties. Machine effectiveness measures such as
20 the percentage of requested fields of view that were focused adequately or the percentage of acquired images that had saturated pixels are measures of how well the invention is setup on and has processed a slide. For a particular slide, the slide suitability score determines the reliability of the other slide
25 scores, and thus, whether those scores should be reported. If a particular slide is anomalous, or if the invention did not operate effectively on the slide, the slide suitability score flags the
30

- 12 -

unacceptable machine condition or slide characteristic so that potentially false results are not used to evaluate the biological specimen slide.

5 The suitability score's primary design goal is to identify the conditions under which the AutoPap® 300 QC had performance limitations during training. The performance requirement of the suitability score for a laboratory for whose processes are within design operation conditions of the automated analysis apparatus is to score no more than 5% of the slides as being unsuitable.

10 Additionally, the suitability score identifies conditions that were not encountered during training. For example, if during training, the stain measures for all slides ranged between 3 and 4, any slide that now measures 2 is outside that range, and as a result, unsuitable. Thus, the range of training data was used to set boundaries for future suitability scoring. Finally, the suitability score identifies slides on which algorithm processing errors were encountered.

20 After the entire slide has been processed, the slide suitability score is determined using features provided by the slide scoring software. The suitability score is archived with other slide scores.

25 The output is an enumeration indicating whether a slide passed or failed the slide suitability tests, and in the case of failure, identifying which test foiled.

30 During the 4x magnification scan the analysis process receives images from the computer 540 FOV boards 568 and sends image information extraction results back to the computer 540 for accumulation.

35 When the 4x scan is completed, the accumulated extracted results 24 are analyzed using the 4x algorithm 18 and the analysis results are passed back

- 13 -

to the computer 540. Computer 540 uses the 4x scan analysis results for 20x scanning prioritization and accumulates the 4x scan analysis results for later slide scoring.

5 Once the 20x scan has been completed, accumulated extracted information is used to perform slide scoring. Slide score is reported and archived.

10 Figure 3 is a slide processing flow chart of the method of the invention to process biological specimen slides. It shows the order in which each analysis sub-system is executed during slide processing. Because processing is distributed in the invention, the low magnification scan 32 and the high magnification scan 38 do not represent a strict sequential process.

15 Those skilled in the art will recognize that the various computers of the invention may execute them in a parallel or serial fashion. During each scan, a number of images are captured and processed, and information about them is accumulated for later use.

20 New slides are loaded under the microscope at step 28 and set up at step 30. During the 4x analysis process 32 the computer 540 extracts information from images during the 4x scan of each slide. The information extracted from each image is used

25 primarily to prioritize image acquisition for the 20x scan 38. Additionally, the information is used for slide scoring 40.

30 The 4x magnification information is used to avoid scanning at high magnification in the following areas of the slide:

- o Blank areas
 - o Areas contained within bubbles
 - o The edges of the coverslip
 - o Areas containing cells that are inadequate for
- 35 analysis at high magnification (for example,

- 14 -

thick sheets of cells).

- Stripe areas, stripes are periodic illumination caused variations appearing on the slide.

The information 34 acquired during the 4x magnification scan is reported as two priority rankings for each 20x FOV: (1) the likelihood that the FOV contains isolated abnormal cells (SIL rankings) and (2) the likelihood that the FOV contains analyzable cellular groups (Group ranking). Additionally, the 4x analysis provides information about the bubble edge content of each 20x FOV (bubble edge flags). The 20x FOV selection analysis 36 uses this information to determine the scan/process sequencing for 20x image acquisition. The scan/process information is passed to system software responsible for acquisition of 20x images.

To provide information for slide scoring, the 4x analysis also derives three histograms - object count, cell count, and cell to object ratio histogram. These histograms provide measures of object populations on the entire slide.

Figure 4 shows the information flow during the 4x scan. The 4x scan is made up of slide scanning, image acquisition, and sequencing 44, 4x information extraction 46 - 48, result accumulation 50, and 20x FOV selection 52. The host processor 540 controls the 4x scanning and image acquisition. For each image acquired, the host processor 540 sequences the image to one of a number of FOV processors 568. The FOV processor 568 extracts information, and results are returned to the host processor 540 for accumulation, 20x FOV selection, and later slide scoring 62 shown in Figure 5.

Still referring to Figure 5, the 20x magnification analysis extracts information 58-60 from

- 15 -

20x scan images. The information extracted by these analysis provides information for slide scoring 62. The primary image extraction analysis is the single cell analysis. The two secondary image extraction
5 analysis are the group analysis and the thick group analysis. A tertiary analysis is the stripe analysis. Information from each analysis is passed to the FOV integration process for collation and transfer to the host processor 540. The host processor accumulates
10 multi-image results. When the 20x scan is completed, the host processor scores the slide.

The stripe analysis detects illumination anomalies caused by inappropriate obstructions in the illumination path, for example, dirt on the coverslip.
15 When stripes are detected, the stripe analysis passes information about the stripes to the FOV integration analysis.

The single cell analysis detects single, non-overlapping nuclei and cytoplasm within 20x images.
20 When such objects are detected, the analysis extracts features for each object and classifies the object as intermediate cell, normal, artifact, or potentially abnormal. The single cell analysis passes object classification, stain evaluation, object feature
25 results, and error information to the FOV integration analysis.

The group analysis detects aggregates of cell nuclei within 20x images. When such objects have been detected, the analysis extracts features for each
30 object and classifies the object as an artifact, possible endocervical, possible metaplastic/parabasal, or possible abnormal group. Since the single cell analysis is executed before the group analysis on each image, information from the single cell analysis is
35 available to the group analysis for object

- 16 -

classification. Object classification, object feature results, and error information are passed from the group analysis to the FOV integration analysis.

5 The thick group analysis detects 3-D aggregates of cell nuclei within 20x images. The difference between thick group and group detection is that the thick group analysis detects aggregates that are so thick that individual nuclei are difficult to distinguish. When thick groups are detected, the
10 analysis extracts features for each object and classifies the object as abnormal or other. Both the single cell analysis and group analysis are executed before the thick group analysis and both can pass information to the thick group analysis for object
15 classification. Object classification results and error information are passed from the thick group analysis to the FOV integration analysis.

The FOV integration analysis takes input from each of the other 20x analysis. All input is
20 integrated, and the instances where different types of objects of interest occur in the same image are computed. For example, when potentially abnormal cells are detected by the single cell analysis, the integration analysis computes how many potentially
25 abnormal cellular aggregates were identified by the group analysis in the same image. The resultant feature is the number of potentially abnormal groups when there were potentially abnormal single cells. All information about the FOV, including features for
30 objects not completely classified (for use by the add-on classifiers), is collated and passed to the host processor 540.

The host processor can classify objects using features from more than one image. Once the required
35 features are available, the host processor completes

- 17 -

object classification for each object and adds the new information to the accumulated slide information. With this information, the slides can be scored.

Each slide is analyzed for abnormality (QC score analysis), composition (Endocervical score and Cell Count analysis), and suitability for machine processing (Suitability score analysis). These analysis are described in Table 1.

Table 1

10	Analysis	Description
15	QC Score	Using information extracted by the 20x analysis from the slide during 20x scan, a score is assigned to the slide that represents the slide's likelihood of containing abnormality. When the c score exceeds a set threshold, the slide should be sent back for human review.
20	Endocervical Score	Using information extracted by the single cell and group analysis from 20x images of the slide, the AutoPap 300 QC system assigns a score to the slide that represents the slide's likelihood of containing endocervical component. When the endocervical score exceeds a set threshold, the slide is considered to contain adequate endocervical component cells.
30	Cell Count Score	Using information extracted by

- 18 -

5 the single cell analysis from 20x images of the slide, the AutoPap 300 QC system assigns a number that represents an estimate of the count of squamous cells on the slide. When this value is above a set threshold, the slide is considered to contain adequate squamous component for analysis.

10 Suitability Score Using information extracted by the 4x analysis and the 20x analysis, an analysis determines whether the slide is suitable for machine processing. If the slide

15 is determined to be unsuitable for scoring, the AutoPap 300 QC system provides a suitability score that represents the reason for unsuitability. When a slide

20 is determined to be unsuitable for machine processing, the only slide score provided to the user will be the suitability score. (All scores are reported to

25 system software regardless of the suitability of the slide. See the Error Handling section.)

Refer now to Figure 6 which shows a system interface schematic diagram. The workstation 542

30 allows users to interact with the invention system. A database 70 of results from each slide run is maintained. The interfaces shown ensure that the following calculations and communications occur:

- 19 -

- The system software 68 makes images available to the analysis on the FOVs.
- The system algorithms 64 process each image and pass results 74 (which includes error conditions if necessary) for that FOV back to the system software 68.
- The system software 68 accumulates FOV results 76 for an entire slide as processing continues.
- The system software 68 provides the analysis with accumulated slide processing results, from which the analysis derive the following scores:
 - QC score 78
 - Endocervical score
 - Cell count score
 - Suitability score
- The algorithms 64 pass the scores back to the system software 68.

- 20 -

Table 2 Algorithm Requirements Allocation

Requirements	Algorithms								
	QC Score	Endoc Score	Suitability Score	Cell Count	4x Alg	Stripe Alg	Single Cell Alg	Group Alg	Thick Group Alg
Suitability for Machine Processing			X						
Endocervical/Transformation Zone Component		X							
Squamous Epithelial Component				X					
Throughput Rate	X	X	X	X	X	X	X	X	X
Specify (Sort Rate)	X								
Cytological Sensitivity	X								

Table 3 Information Analysis and Supporting Algorithm Dependencies

Information Analysis Algorithms	Supporting Algorithms					
	4x Alg.	20x FOV Selection Alg.	Stripe Alg.	Single Cell Alg	Group Alg	Thick Group Alg
QC Score		X		X	X	X
Endocervical Score		X		X	X	
Suitability Score	X		X	X	X	X
Cell Count				X		
Add-On Classifiers				X	X	N/A
20x FOV Selection Alg	X	N/A				

- 22 -

Each of the information analysis operations (QC score, endocervical score, suitability score, and cell count) is dependent upon supporting analysis. Table 3 shows these dependencies. The horizontal axis lists the supporting analysis. The vertical axis lists the information analysis. (Note that Add-on Classifiers and 20x FOV Selection analysis are included in this list because they require information from supporting analysis. The Add-on Classifiers and 20x FOV Selection analysis in turn provide information to the information analysis, so they appear in both the horizontal and vertical axes.) An x indicates that an information analysis is dependent upon a supporting analysis.

For example, the Specificity requirement is solely dependent upon the QC score analysis, so the specificity requirements must be met by QC score analysis. The QC score analysis, however, is dependent on five information extracting or supporting analysis, see Table 3.

Refer now to Figure 7 which shows the apparatus of the invention during training. Computer training programs generate the classification processes of the apparatus of the invention using a representative set of examples. Each example comprises an object with a known diagnosis. Each object is represented by a set of discriminative features. An object can be a single cell, a group or a thick group of cells, or a slide.

The machine training process must deal with pap smear qualities that present major challenges. These qualities include large variations caused by system bias, random noise, staining, sampling, fixation, slides, cells, and so on. To cover all possible variations, machine training requires sufficient training cases. Also, it important to derive robust

- 23 -

features, which are less sensitive to the variations, and to carefully validate the decision structures and selected features.

5 Abnormal slides may have over 100,000 normal cells and as few as 10 abnormal cells. Detecting these slides requires high sensitivity to abnormal cells, yet demands an extremely low false-positive rate for normal objects.

10 Machine training includes the three phases shown in Figure 7, the preparation phase 80, object training phase 82 and slide investigation phase 84. The preparation phase 80 involves, development of the machine training strategy, development of training tools, and the development of features. The apparatus
15 of the invention provides a tool to understand the clinical significance of slide and object features, develop the protocol for determining diagnostic truth, define object classification structures, identify clinically useful features for each cell
20 classification decision, categorize training objects and determining sample sizes, characterize objects of interest by prevalence, and define performance goals.

Refer now to Figure 8 which shows a guide for classification architecture that specifies
25 discriminant features, performance goals, and object categorization. Figure 8 shows the components of a representative classification architecture for objects 318. Clinically meaningful, discriminant cell features and classification performance goals 320,
30 322, 324, 326 are used to define the classification architecture. The architecture is organized as a decision tree. The inputs and outputs, as well as discriminant features and performance goals, are specified for each node of the tree. Outputs for
35 object classification include classifications 328,

- 24 -

330, 332 and 334. Ideally, the classification architecture may be used as a basis for developing an object classification process.

5 The inputs to the preparation phase 80 are the results of an algorithm clinical knowledge integration stage. The objective of this stage is to clearly translate clinical knowledge and guidelines into detailed classifier designs and implementation plans. The tasks are to define or develop a machine
10 architecture, a strategy for feature design, implementation, and testing, a strategy for algorithm evaluation and refinement, an image collection strategy, a strategy and protocol for database management, a work flow and protocol for data
15 collection, and a quality assurance strategy for machine training and implementation.

Refer now to Figure 9 which shows a data flow diagram of one embodiment of the processing steps in machine training. After the development of the
20 machine architecture and training strategies, a data flow for machine training can be determined. To satisfy specify machine training requirements, features and tools are developed or modified.

Training slides 96 are acquired and cells of
25 interest are identified 98. Then, images containing the cells of interest are acquired 100, and the objects are segmented 102. The features of the segmented objects are extracted 104 and the objects are labeled. After a sufficient number of object
30 features and their corresponding labels are collected, features that allow discrimination between classes of objects are selected 106. The selected features are used as the basis for the construction 108 of the object classifiers. The object classifiers are
35 evaluated 110 to determine if the resulting

- 25 -

classifiers sufficiently meet the processing requirements 116. All data and results are stored and manipulated in a common database management system 112. The database system 112 also provides a platform to support data visualization.

During the algorithm clinical knowledge integration stage, cytotechnologists or cytopathologists define discriminate features. Algorithm developers evaluate the defined discriminant features and define the corresponding computer measurements. In addition, the algorithm developers research literature on automated cytology for information on other computer features that are known to be useful for automated cell classifications. The current imaging system capabilities are evaluated to further refine the feature list. A comprehensive feature list is then generated based on the gathered information.

The selected features are defined in detail and their implementation strategies are specified. After the completion of the specification, the system software group members can implement the features; the algorithm group members verify the feature implementation by comparing the results of independent implementations. After verification, the features are used as the basis for machine training.

Each component of the machine training data flow requires a set of tools to facilitate the successful implementation of the specific training tasks. Two classes of tools are required: object acquisition tools and classifier development tools.

The object acquisition tools include software and hardware tools for image acquisition, image processing and feature extraction, and object labeling. The tools developed include the following: slide bar code reader

- 26 -

and PC-based cell coordinate acquisition station, the semi-automatic cell capture program, image saving mechanism for automated acquisition, image processing and feature extraction programs hosted by the
5 algorithm development station, the object labeling stations for entering the object class labels to the database, and the object labeling program.

After training data is collected, many pattern recognition tools are developed to facilitate the
10 classifier development work. These tools include feature selection tools, feature transformation tools, classifier building tools, and classifier valuation tools.

After the preparation phase 80, the algorithm
15 development process moves to the object training phase 82. The inputs to phase 82 include the developed features and training tools. The outputs of object training phase 82 are the developed object classifiers, which have satisfactory performance at
20 the object level. The object training process is shown in Figure 10. The process follows these steps. Images of the cells of interest are semi-automatically acquired 118. The acquired images are used to refine the segmentation processing 120. The segmented
25 objects are labeled with their true classes 122. The labeled objects are used to perform object classifier training 124. An object classifier is implemented in the system and tested. If the performance of the classifier is unsatisfactory, then objects that the
30 classifier failed to label are acquired in an automated object acquisition mode 128. Steps 122 through 128 are repeated until the performance of the classifier is satisfactory. After a satisfactory classifier is constructed, it may be productized 126
35 and used to classify biological specimens on

- 27 -

microscope slides.

The semi-automatic image acquisition process is shown in Figure 11. This process comprises three basic stages: selection, acquisition, and quality controls

5 132. A slide is selected from the training slide set and qualified 130 subject to quality controls 132.

Training objects are selected and qualified 134. A slide level diagnosis is established 140. Training object coordinates are acquired 136. Training object

10 images are also acquired 138 from machines 142.

The acquired images are processed by an image segmentation algorithm. The baseline image segmentation algorithm is designed to be as insensitive to image quality variations as possible.

15 This baseline segmentation algorithm can be modified as desired to match the specific image quality and application scenarios. The segmentation refinement is an iterative process through multiple training and testing cycles. The refined segmentation process can
20 be fast ported to speed up the remaining tasks of the training process.

The refined image segmentation process is applied to the acquired images, and the features corresponding to the segmented object masks are extracted and stored
25 along with the object masks. The original images, the segmented object masks, and the features are used by a cell labeling program to label and verify the objects. The object labeling task is conducted by a cytotechnologists following predetermined standards.

30 After a sufficient number of the objects are acquired and labeled, the training process moves to the object classifier training stage. Object classifier training uses approaches including the development of multi-stage classifiers following a
35 redefined structure, and the balancing of training

- 28 -

data for labs, machines 142, and object types.

As shown in Figure 12, the acquired images from the semi-automatic cell library 154 are used to generate the initial classifier 146. The initial
5 classifier 146 is implemented in the system, and an automated data collection process is conducted to acquire additional training samples. The automated collection data and the semi-automatically acquired data are used to train the subsequent classifiers.

10 The automated data collection processes 156, 158, and 160 use a focus of attention method in which the newly acquired data are the samples that the current classifiers 148, 150, 152 failed to classify correctly. In addition, this training method requires
15 that the implementation of the intermediate classifiers be the basis for further data collection. As a result, sampling errors can be corrected.

It is not possible to train processes on paper and wait until the whole process is developed before
20 implementing the classifiers. This unique property makes it difficult to follow strictly the conventional waterfall process of software development. Instead, the process group follows a spiral method of development, evaluation, and planning. To collect
25 training data, the same set of slides can be processed multiple times to capture the system variations and improve the robustness of the resulting classifiers. To assure the integrity of the training process, both the machine 142 and the data have to be qualified as
30 part of the process. Many machine monitoring and failure recovery activities are applied throughout the training process.

The object classifier training process is controlled by a classifier evaluation process, which
35 comprises interactive data visualization, feature

- 29 -

evaluation, error rate estimation, and clinical verifications. Processes are adjusted and expanded until the performance goal is achieved or failure causes are accepted as system limitations.

5 As part of the object classifier development process, a classifier development strategy is applied at each stage of the training. The classifier development strategy is shown in Figure 13. A preliminary classifier is developed using manual or
10 semi-automated data from a data collection process 162 as described above. Once a preliminary classifier 164 is developed, the classifier is evaluated 166. The evaluation process comprises statistical analysis and interactive data display and visualization methods to
15 characterize the data distribution and decision boundaries. The basic steps of the classifier evaluation process include the following checks. The error rate of the classifier is determined and assessed 188. The features and images of objects that
20 failed in the classification process are examined and the causes of the failures are determined 186. The features selected by the classifier and the rules of the feature applications are compared with the expected clinically useful features and how they are
25 used 190. The process performance is compared with the predetermined goals 190. If the evaluation results show satisfactory performance 168, the classifier development process is completed 170. If the evaluation results determine the current
30 performance level falls short of the desirable goal, additional decision logic is applied 172.

 The next decision is whether an additional classifier should be built 172. This decision depends on the current classifier's sensitivity to the objects
35 of interest. If the sensitivity is high enough (for

- 30 -

example, >60%), the process switches to the development of an additional classifier 176, and the evaluation process is activated after the next classifier is developed.

5 If the sensitivity is low (for example, <50%), it is difficult to develop any additional classifiers. At this point, step 178, the classifier can be accepted in spite of its limitations 180, or the classifier can be further refined 182. If the classifier's
10 performance level is not acceptable, and possible improvements to the classifier have been identified, the current classifier may be refined. The refining process comprises acquiring additional data 184 and then rebuilding the classifier using the enlarged data
15 set and refined features.

Process accuracy has the highest priority among performance goals. However, to facilitate the productization of the developed processes, process robustness, throughput, and implementation maturity
20 and maintainability are also evaluated, and if necessary, improved in the process development process.

After the developed object classifiers are integrated into the system, a slide training readiness
25 test is conducted to ensure the quality of object training phase. If the slide training readiness test is completed and the results are satisfactory, the object classifier training phase is considered complete, and the development process moves to the
30 slide evaluation phase.

Quality assurance processes are applied throughout the object training phase. Quality control procedures for data collection, object labeling, and classifier development are defined. Most of the
35 procedures include auditing of the data, comparing the

- 31 -

results with the pre-defined goals, and performing discrepancy analysis.

Now referring to Figure 14, the slide investigation phase 84, is illustrated. This is the
5 last phase in process development. When the slide investigation phase is completed the developed process is specified and then implemented into the machine to analyze biological specimen slides. The inputs to
10 this phase are the successfully developed object classifiers 192. The objectives of this phase are to further evaluate the object classification results. This evaluation is based on a large set of new slides 194 and the effectiveness of the accumulated classification results on a slide by slide basis.

15 As shown in Figure 14, new slides 194 are acquired and processed. Prior to this, an investigation protocol is developed that defines the type and selection criteria of the slides, the configuration of the slide processing, and the data
20 formats. As part of slide processing 196, images having significant objects detected by the object classifiers are saved, and the object classification results are accumulated for each slide. The results are used to evaluate the object classification
25 results.

To specify the truth of the detected objects, cytotechnologists label the images having significant objects in an object analysis procedure 198. The truth labels are used to determine the accuracy of the
30 object classifier. For slides having a high prevalence of abnormal cells and a low number of detected objects of interest, a sensitivity analysis is conducted that saves and labels all the images processed by the process. The sensitivity analysis results also help
35 determine the accuracy of the object classifier.

- 32 -

The object classification results are accumulated over the whole slide, and a slide analysis 202 is conducted. Based on only these results, the accuracy of slide classification is determined and compared with the desirable process performance goals. The object classification results are the foundation of the slide scoring processes; therefore, if the results are very unsatisfactory, a detailed analysis may be required to determine the feasibility of the current system and process approach.

The accumulated object results are used to derive slide features 200. Additional slide features are extracted based on the integration of the results from multiple object classifiers. The slide features and the diagnosis of the slides are used to develop slide scoring processes.

The slide scoring processes 204 are developed using a training process similar to the object training processes. The developed slide scoring processes are then evaluated to validate the overall development process.

In addition to the accuracy analysis, the repeatability and throughput of the object accumulated results are also evaluated to determine whether the object classifiers provide a solid foundation for an acceptable product.

Process robustness, throughput, and implementation maturity and maintainability are also evaluated, and if necessary, improved in the development process.

After the developed slide scoring processes are integrated into the system, a slide scoring evaluation test 206 is conducted. If the results are satisfactory, the slide evaluation phase is complete.

Quality assurance processes are applied

- 33 -

throughout the slide evaluation phase. Slide diagnostic truth is establishment follows predetermined guidelines. System qualification and integrity are established during slide processing. In
5 addition, verification of the code is done.

To illustrate the operation of the invention the details of endocervical group analysis will be described. Endocervical group analysis is performed on two sources of input. For each slide, it receives
10 20x fields of view from the image acquisition module. These FOVs are images that a 4x magnification analysis has determined to have a likelihood of containing groups of cells, whether normal endocervical or potentially abnormal cell groups. During endocervical
15 group detection normal endocervical and abnormal cell groups contained in an image are also detected. The invention classifies all segmented objects in each image and further processes the classification results with a 20x FOV integration analysis for slide
20 classification. Whole image results are also used during thick group analysis.

The endocervical group analysis of the invention starts with images of a cytological specimen and cell results using whole image features. The endocervical
25 group analysis outputs whole image features, which it passes to the thick group process, object and group classification data, if cell groups are found and sent to 20x FOV at integration step and feature active flags.

30 The endocervical group analysis proceeds initially as follows:

- o Results from 4x magnification processing and checks bubble edge locations to ensure that all areas inside bubbles are excluded from
35 processing. Areas within bubbles are eliminated

- 34 -

from processing because presentation of cells within bubbles may not be reliable.

- o Cover slip edges are detected and that all areas outside of the area bounded by cover slip edges are excluded from image processing.

- o Image acquisition assures that images passed to the group process for processing conform to image quality specifications. System software assures that images are acquired based on rules provided in 20x FOV. Endocervical group analysis uses nuclear detection, grouping, and classification techniques to identify probable groups of endocervical cells. The presence of normal endocervical cells in a specimen is a partial indication that the sample was taken from the proper anatomical location, i.e. at the squamo-columnar junction, which is the transformation zone of the uterine cervix. Groups of cells that lie with densely overlapping nuclei are not detected by the endocervical group process. These kinds of groups are detected and classified by thick group analysis. Free lying cells - cells having isolated nuclei - are detected by the single cell processing.

The endocervical group analysis functions to provide an endocervical score which helps to determine whether a specimen has been adequately sampled by the physician who originated the slide. In addition supporting information is supplied to generate an anomaly score. Group analysis operations are divided into five major steps. Processing for each image is completed before the next image is sequenced to the FOV board for processing. Results of each processed field-of-view are accumulated for slide scoring.

Step 1: Once an image, which in this context is

- 35 -

equivalent to one 20x FOV, has been made available for endocervical group processing, features are calculated in preparation for Step 2 classification. This step is called whole image feature extraction.

5

Step 2: The process uses the whole image features provided by the single cell process and calculated during Step 1 to classify each image into one of two categories: "no potential for containing identifiable cell groups", or "may contain identifiable cell groups". If an image is classified as having no potential for containing identifiable cell groups, processing ceases for that FOV and results are reported to the 20x FOV integration process and to the thick group process. This step is called FOV classification.

10

15

Step 3: For each FOV that passes Step 2, the process segments objects that may be groups of cell nuclei. If no objects are identified in an image, processing stops for that FOV and results are reported to the 20x FOV integration module and to the thick group process. This step is called image segmentation.

20

25

Step 4: Next, the process measures feature values -- such as size, shape, density, context and texture -- for each potential cell grouping detected during Step 3. This step is called feature extraction. Features for each object are passed to Step 5.

30

- 36 -

Step 5: Finally, the process classifies each detected object in an FOV using the feature values obtained in Step 4. Processing stops here and results for each object are reported to the 20x FOV integration process. This step is called object classification. Classification rules are defined and derived during machine training.

Processing occurs sequentially, one FOV at a time. Each FOV is passed through the five steps outlined above until a determination is made about it. An individual FOV is always passed from Step 1 (whole image feature extraction) to Step 2 (FOV classification). If a determination can be made at this point that the image could not contain potential cell groups, it is eliminated from further processing and its results are passed to the 20x FOV integration process. Otherwise, it is passed to Step 3 for image segmentation. Step 3 either eliminates the image from further processing or decides that it contains segmentable objects which may be cell groups, in which case the FOV is passed to Step 4 for feature extraction. The image is then passed to Step 5 for object classification. Accordingly, there are three possible outcomes for each FOV that enters the flow:

1. Could not contain identifiable cell groups as determined by Step 2;
2. Does not contain identifiable cell groups as determined by Step 3; or
3. Contains objects which are identified as normal endocervical cell groups, potential abnormal cell groups, metaplastic/parabasal groups, uncertain endocervical or metaplastic/parabasal groups, or artifacts.

- 37 -

The first step of processing is whole image feature extraction. A whole image is equivalent to one 20x FOV. The goal of the first step is to measure general image properties in an FOV that may indicate the presence or absence of potential cell groups. This information is then communicated to Step 2, FOV classification. No FOVs are eliminated by Step 1. They are all passed to Step 2 for classification.

Features at this level of processing consist of measurements that can be taken from a whole image. These are features that tend to indicate image content, texture, and contrast across the whole image. Three whole image features are received from the single cell process. These are the single cell results, as input to the endocervical group process. In addition, Step 1 of the endocervical group process calculates two additional whole image features. Generally, the single cell whole image features are measuring the global image intensity and image content of an FOV, while the endocervical group features are measuring texture and contrast in an image that may be the result of edges of potential objects in an image.

Whole image features received from the single cell processing are:

1. High count,
2. High mean, and
3. Low threshold.

The endocervical group process calculates two additional features that measure texture and contrast that may be the result of edges of potential objects in an image. These features are:

4. 2x1 edge magnitude, and
5. 5x5 dark edge strength.

- 38 -

Image (or FOV) classification is Step 2 of the endocervical group process's process. Using the five whole image feature values that were described above, the process uses three box filters to test each image for potential of containing identifiable endocervical or other cell groups. If an image passes these three boxes, processing continues to the next step. If an image does not pass, processing stops for that FOV and results are passed to the thick group process and to the 20x FOV integration process.

Image segmentation is the process of identifying objects of interest within a gray scale image and creating an object mask, which is a binary image that represents the objects of interest. Each area of interest is represented as active pixels in the object mask. The group process's strategy for locating potential endocervical groups is to identify probable cell nuclei and determine whether any nuclei lie in close proximity in sufficient numbers to be part of an endocervical group. Once a group has been segmented, nuclei and cluster masks are provided for each object for use by the feature extraction software described herein.

Since endocervical cell nuclei tend to lie in such extreme proximity as to be overlapping, and since a Papanicolaou-prepared cervical smear typically contains a large number of objects that mimic cell nuclei in appearance, most of the image processing effort is expended to identify cell nuclei. Image segmentation for the group process follow four major steps. Initially, objects other than nuclei are removed from an image. Next thresholds are calculated to identify each nucleus. Segmented objects are checked to ensure that they represent nuclei. Finally, nuclei are examined to determine if they lie

- 39 -

in clusters.

Nuclear segmentation progresses through four major steps including background object removal, image thresholding, object refinement, and nuclei clustering. These steps contain several substeps. Background object removal comprises two steps:

1. Large object removal, and
2. Small object and odd shaped object removal.

Image thresholding comprises:

3. Threshold calculation.

Object refinement comprises three steps:

4. Weak edged object removal,
5. Nuclear separation (for nuclei whose masks are joined), and
6. Small hole filling.

Since nuclei have a finite size range, all objects outside that range are removed first. Along with small objects, thin objects are removed since nuclei tend to be round or ovoid. A threshold is calculated for the identification of areas of the image which may contain nuclei. Once areas for nuclei have been identified, each potential nucleus is checked to ensure that it has sufficient edge contrast around its entire perimeter. Those that do not meet this criterion are removed. Nuclei masks are checked to see if they may actually be two nuclei connected together. If so, they are separated. Finally, if small holes remain in the nuclei masks, they are filled.

Large objects may be cell cytoplasm, mucous streams, or non-cellular artifacts. Such objects are removed from further consideration by an image closing operation that uses a structuring element larger than the largest nucleus. This closed image is then

- 40 -

conditionally eroded using the original image as the conditional image. The result is an image with no detail smaller than the structuring element used for the closing operation. The original image is
5 subtracted from the closed image, yielding an image that contains only the detail of the largest possible nuclei and smaller objects. The product of this step is called a residue image.

Objects that are too small or thin to be
10 endocervical nuclei are removed by checking to see if any of the residue objects in the residue image can be removed by opening with a structuring element smaller than the smallest acceptable nucleus. Essentially, this check makes sure that each object's minimum
15 diameter is larger than a minimum size. If an object does not pass, it is removed from the residue image.

During the development of the nucleus segmentation process, two different methods of thresholding nuclei were found to be effective. Each
20 had its own advantages. One method segmented the majority of nuclei and generally segmented nuclei entirely with little oversegmentation. The problem was that it created a large number of false nuclei as well. The second method led primarily to segmentation
25 of nuclei of interest with few false nuclei. This method, however, did not do an adequate job of segmenting each nucleus in its entirety. It was therefore decided to use a combination of the two segmentation methods. The first method provides a
30 formation mask into which a nuclear seed mask - the result of the second method - can be grown. The net result is well formed nuclei with few false nuclear segmentations.

Each thresholding method creates an image. These
35 two images are used to apply thresholds to a modified

- 41 -

residue image. Thresholding occurs pixel by pixel. Information about each pixel and its environment is used to decide whether a pixel is part of the background or part of a nucleus. The context of each pixel is found in the pixel itself, the pixel's immediate surroundings, and the entire image around the pixel. Each pixel's threshold value is calculated from individual pixel values, local area parameters, and whole image parameters. The formation mask uses local area information for thresholding while the nuclear seed mask uses all three.

The original image provides background information needed for the thresholding of the formation mask. If nuclei are free lying, their thresholds should be higher than if the nuclei lie in dark backgrounds. The original image is adjusted so the background information is at an appropriate level for thresholding the modified residue image. The image is adjusted by scaling, biasing, blurring and clipping. Scaling and biasing ensure appropriate threshold levels. Blurring removes information other than background from the threshold image. Clipping assures that nuclei have at least a minimum residue strength before they are included in the mask.

Several images are combined to form the threshold image for the nuclear seed mask. Additional object information is added to the modified residue image for increased effectiveness. Images are created to reflect information in this new modified residue (nmr) image. They are combined and used as the threshold image for the nuclear seed mask. Additional object information is added to the modified residue image using binomial filter residue near objects that have been closed. This filter residue is an additional measure of contrast in areas around objects and

- 42 -

provides a more robust selection of objects.

Two whole image parameters - edge and brightness - are used to create a threshold image for the nuclear seed mask. Edge, which is a measure of medium to high spatial frequency content of an image, is the average of a binomial filter residue and a dilation residue of the original image. Brightness is the average pixel value in the original image.

Three additional images are created from the nmr image. The first is an object edge image. It is a dilation of the nmr image with the nmr and an eroded nmr subtracted from it. This is expressed as:

$$(NMR (3 \times 3)) - NMR - (NMR \ominus (3 \times 3)).$$

The result is an image which has object edges only. The second image is the nmr image for selected objects. Each object has some strength associated with it. If that strength is greater than a value calculated from the edge and brightness values just described, the object is part of the second image. This image increases the threshold for those objects that may be two nuclei close together. The increased threshold helps to ensure that nuclei are segmented separately. The third image is just the new modified residue image. Each of the three images is scaled and added together to form the threshold image for the nuclear seed mask. A comparison is then made between the combined image and the nmr image. Resulting nuclei are conditionally dilated with the formation image.

There may be objects other than nuclei represented by the nuclei mask. One property that can be exploited to remove false nuclei is that nuclei are circular objects that are generally darker than their surroundings. A circular edge that goes from bright to dark as one moves from background into the nucleus

- 43 -

should exist near the perimeter of each object. Two image parameters, edge and brightness, determine how strong edge contrast needs to be before it is considered to be sufficient. Edge strength is checked for each object in four directions: vertical, horizontal, diagonal up and diagonal down. If edge contrast is strong enough in all four directions, the nucleus is retained. Otherwise the object is removed from the nuclei mask.

Objects may remain as part of the nuclear mask that actually represent two nuclei. If the objects can be separated based on shape, that is if the mask object has an hourglass shape, they are separated. Each object is eroded and checked to see if it breaks into two reasonably sized objects. If so, the original object is split along a line that goes between the two objects in the eroded version of the object. If an object did not split during the erosion process, it is eroded again then the object is checked to see if it split. If the object does not split after three attempts, or if the object gets too small, the object is considered to be one nucleus and no modification is made to the nuclear mask.

Nuclei that remain in the nuclei mask at this stage may need to be cleaned up. Nuclei may have small holes in them or jagged edges on their perimeters. These flaws are corrected by closing and opening the nuclear mask where appropriate.

Once nuclei have been recognized, nuclear clusters need to be identified for use in the feature extraction step. Groups of endocervical cell nuclei tend to lie in close proximity to each other. Just how close they lie tends to depend on nucleus size. In other words, when nuclei are large, the distance between nuclei in endocervical groups is large. Small

- 44 -

endocervical nuclei lie close together within a group.

To achieve clustering dependent on nuclear size, a nuclear mask is separated into four parts. The four masks contain nuclei that remained after successively
5 larger erosions. This way nuclei are ordered by size. The mask that holds the largest nucleus is then dilated. The mask that holds the second largest nuclei is dilated, but by a smaller amount. Each of the four masks is dilated by an amount related to the
10 size of nuclei contained in the mask. All four masks are logically OR'd together. Nuclei that are part of the same group will now overlap, forming clusters. Objects or clusters that are larger than some minimum are retained as potential endocervical groups. Nuclei
15 associated with the retained clusters are retained, and both the nuclei mask and cluster mask are made available for use by the feature extraction step.

After segmentation is complete, segmented objects are passed to the group process's feature extraction
20 level. Once possible cell groups are identified by image processing, features (which are properties of segmented objects) of each cell group are measured. Specific features are chosen which provide the best discrimination between different cell group types and
25 artifact groups.

A wide range of feature types is necessary for good identification of segmented groups. Features used can be looked at and categorized in two different ways. They can be identified by the kind of
30 information they measure - object shape, size, and texture for example. Or, features can be identified by what part of an image they measure - the object of interest, a small area around the object, the whole image, or combinations of these features. Object area
35 measurements attempt to characterize the object

- 45 -

itself, while other area measures try to characterize the object's surroundings; they are an attempt to provide cell group contextual information. Combined features measure how an object appears compared to its surroundings.

Features are grouped for discussion first by the image area they are calculated over and second, by what type of measurement is being made. Features are uniquely named and identified logically in the form of feat 099, which means, of course, Feature 99. See Appendix A for information about the entire group of features and how they are allocated to the various cell classification levels. After features have been extracted for an object, object classification is finished.

Features measured in the cluster area form the most diverse set. They may be calculated from nuclei, cytoplasm density, size, shape, texture, orientation, or from a ratio of nuclear and cytoplasm density, texture, or area. Features used from the cluster area only are described next. They are categorized according to the general kind of information they measure.

Feature 7 is the average intensity - or pixel value (0 to 255) - of the non-nuclear cluster area. This feature provides information on whether a cluster is located on top of other objects or if it's densely packed.

Feature 111 is the average of nuclear edge contrasts. If a nucleus is oddly shaped, it may not have a very reliable measure; in other words it may be poorly segmented. Only so-called "good" nuclei are used in conjunction with this feature. Feature 111 is the average difference between pixel intensity values two pixels inside good nuclei, and pixel intensity two

- 46 -

pixels outside good nuclei. This gives some measure of nuclear intensity without influence from background light levels. It may also provide a measure of segmentation reliability.

5 Feature 116 is the standard deviation of the average optical density of "good" nuclei. This measure is an indication of how different one nucleus is from another within the group; that is, how much nuclear intensity varies from nucleus to nucleus.
10 This measurement is taken only on "good" nuclei.

 Feature 119 is the average of the normalized nuclear optical densities of good nuclei. Optical density of each nucleus is normalized by subtracting the average optical density of a ring around nuclei
15 from the average optical density of the nuclei. If objects other than nuclei themselves are in the optical path, normalization removes the effect they have on measuring nuclear optical density.

 Feature 5 is the area, in pixels, of the cluster
20 mask.

 Feature 103 is the average area, in pixels, of "good" nuclei.

 Feature 104 is the standard deviation of the areas of "good" nuclei.

25 Feature 125 is the average perimeter of all nuclei. The perimeter of each nucleus is the average of the perimeter of the nuclear mask and the perimeter of the same mask expanded by one pixel.

 Feature 126 is the standard deviation of the
30 perimeters of all nuclei.

 Feature 6 measures the compactness of the cluster mask. Compactness is calculated as the perimeter squared divided by the area of the object. The compactness of a circle is about 1.27; the compactness
35 of a square is about 16. This is a measure of how

- 47 -

smooth and/or round the perimeter of an object is. The higher the number the less smooth the perimeter or round the object.

5 Feature 99 is the number of good nuclei in the group. "Good" is defined as having a compactness of less than 12. When compared with the number of nuclei segmented, this feature can give some indication of how well segmentation worked. It may also be interpreted to give an indication of how difficult the
10 segmentation task was.

Feature 106 is deviation of the nuclear compactness of all nuclei. This measure can indicate how much the shapes of nuclei in the cluster vary.

15 Feature 8 is standard deviation of the pixel intensity of non-nuclear pixels in the cluster. This gives some indication of the texture in the cytoplasm around the nuclei.

Feature 13 is called 2x1 (2 pixels horizontally by 1 pixel vertically) bright edge directional
20 disparity within the cluster. Bright edge directional disparity is a combination of two measures. Pixels are examined to find those that have darker neighbors on both sides horizontally, then those that have darker neighbors vertically. For each pixel that
25 passes the neighbor test, the magnitude of the difference is recorded. The magnitude of differences for all horizontal pixels are summed. Then all pixels in the vertical are summed.

Feature 13 is calculated as the minimum of these
30 two values divided by the sum of the two. It provides a measure of whether there are significantly more relatively bright pixels in one direction versus the other. This feature shows whether there is some directionally dominant texture in the cluster. In
35 this case, the texture is very fine, or of a high

- 48 -

spatial frequency.

Feature 14 is called 5x1 bright edge magnitude. As with the directional disparity described in Feature 13, this measure is made up of two directions:
5 horizontal and vertical. In this case, rather than looking for pixels that are surrounded by dark pixels in one direction, groups of three pixels are examined to see if they are bounded by dark pixels on both sides horizontally and vertically. Feature 14 is
10 calculated by squaring the two measures, summing them, then taking their square root. This feature gives a measure of how much edge there is in the cluster that is about three pixels wide in either direction. It also gives a measure for the amount of texture there
15 is that has bright spots about three pixels in size.

Feature 15 is called 5x1 bright edge directional disparity. This measure is the same as Feature 13 except that groups of three pixels are checked for dark neighbors rather than a single pixel. This
20 feature measures directional dominance of texture in the cluster. The texture measured here is bright regions three pixels wide or high.

Feature 18 is 3x3 bright edge strength. The process searches for pixels that have dark pixels
25 around them in all directions. The difference between the bright pixel and its surrounding pixels are accumulated for all such pixels in the cluster. The accumulated figure is normalized by the total number of pixels in the cluster. This feature measures the
30 amount of texture in each cluster that comprises bright regions about one pixel in size that are surrounded by darker pixels on all sides.

Feature 20 is 17x17 bright edge strength. This is similar to Feature 18 except that the bright region
35 searched for is about 15x15 pixels rather than about

- 49 -

one pixel. This measure looks for texture with bright spots that measure 7 to 10 microns on a side.

Feature 21 is the magnitude of the 2x1 dark edge. This feature is the same as Feature 14 except that
5 dark pixels are searched for rather than bright regions 3 pixels wide. This measure is of the total amount of dark area covered by single pixels bounded in two directions by bright area.

Feature 25 is the magnitude of the 9x1 dark edge.
10 This feature is the same as Feature 21 except that regions of 7 pixels in width or height are searched for that have bright neighbors.

Feature 27 is 5x5 dark edge strength. This feature is similar to Feature 18 except that a dark
15 region is searched for that is about 3x3 pixels rather than about one pixel. This measure looks for texture with dark spots that measure 1 to 2 microns on a side.

Feature 28 is 9x9 dark edge strength. This is the same as Feature 27 except the size of the dark
20 region searched for is about 7x7. The texture this feature measures are dark spots about 4 microns on a side.

Feature 29 is 17x17 dark edge strength. This feature is the same as Feature 20 except that dark
25 regions are searched for rather than bright.

Feature 30 is called mean 3x3 blur residue in the cluster. The process measures the absolute difference between a 3x3 binomial filtered image and its original. The average pixel value of this difference
30 is Feature 30. This feature measures high spatial frequency in the cluster.

Feature 31 is called standard deviation of the 3x3 blur residue in the cluster. This measure gives some indication of how uniform high spatial
35 frequencies are within the cluster.

- 50 -

Feature 121 is the average of the standard deviation of pixel values within each of the "good" nuclei. This gives an average range of values for pixels in the cluster.

5 Feature 136 is 3x3 dark edge strength within the "good" nuclei. This measure is the same as that described in Feature 18 except that the search is for dark regions within the "good" nuclei. This feature measures small dark spots in nuclei.

10 Feature 9 is average internuclear skiz distance. Skiz points are points that are not within nuclei and are equidistant from their two closest nuclei. Feature 9 is the average distance from the nuclei of each of these points in the cluster. This measure
15 gives a value for how closely oriented nuclei are within the cluster.

Feature 42 is polarity area within the cluster. This is a measure of how many ovoid dark areas are within the cluster area. If there are a lot of dark
20 regions that are ovoid, the number here will be large. If dark areas tend to be circular, the number will be smaller.

Feature 101 is the average minimum nucleus-to-nucleus distance between the "good" nuclei. This
25 gives a value for how closely oriented "good" nuclei are within the cluster.

Feature 2 is a measure of the nuclear/cytoplasm relationship. It is the ratio of the standard deviation of the average nuclear intensities to the
30 standard deviation of the non-nuclear cluster intensities. This measure indicates how much average intensities of nuclei vary when compared with that of surrounding cytoplasm.

Feature 3 is also a measure of the
35 nuclear/cytoplasm relationship. It is the ratio of

- 51 -

the total nuclear area to the total non-nuclear area. This measure gives information about the nuclear/cytoplasm area ratio for objects in the cluster. It also provides an indication of the
5 proximity of nuclei within the cluster.

Contextual information can be thought of as what company each object keeps, or what's around the edges of an object. One way to get a measure of contextual information is to calculate features in regions
10 external to the cluster. Several features relating contextual information of this kind are used in the object classifiers. They are described next.

Feature 77 is the magnitude of the 2x1 dark edge in the ring around the cluster. This is the same as
15 Feature 21 except that it's calculated in the ring around the cluster rather than in the cluster.

Feature 81 is the magnitude of the 9x1 dark edge in the ring around the cluster. This is the same as
20 Feature 25 except that it's calculated in the ring around the cluster rather than in the cluster.

Feature 83 is called 5x5 dark edge strength. This is the same as Feature 27 except that it's calculated in the ring around the cluster rather than
in the cluster.

25 Feature 84 is called 9x9 dark edge strength. This is the same as Feature 28 except that it's calculated in the ring around the cluster rather than in the cluster.

Feature 86 is the 3x3 blur residue mean. This is
30 the same as Feature 30 except that it's calculated in the ring around the cluster rather than in the cluster.

Feature 88 is 3x3 blur residue skewness. This measure indicates how uniform high spatial frequencies
35 are within the ring around the cluster.

- 52 -

Feature 91 is called 7x7 blur residue standard deviation. The process measures the absolute difference between an original image and the image blurred with a 7x7 binomial filter. Feature 91's value is the standard deviation of the residue from all the pixels in the ring around the cluster. This feature measures uniformity of the residue values from the 7x7 blur. Blur uniformity indicates how variable texture is within the area being measured.

Feature 72 is the maximum polarity in the ring around the cluster. This feature is the same as Feature 42, except that its polarity is broken down further into its maximum direction, its direction perpendicular to the maximum, and its direction to 45 degrees on either side of the maximum. Polarity in the maximum direction in the ring around the cluster gives the value for Feature 72.

Combined features are calculated using both cluster and ring measures. These features provide another measure of contextual information about a cluster. First, calculate feature values for both the cluster and a ring around it. Next, combine that information in a way that reduces the effect that background "stuff" has on the cluster itself. These so-called "normalized" features are described in this section.

Feature 48 is the average brightness of the normalized cluster. The average intensity of the cluster is divided by the average intensity of the ring around the cluster. This ratio is average brightness.

Feature 51 is the magnitude of the 5x1 dark edge. This feature is similar to Feature 21 except it searches for a 3-pixel region of dark area rather than a 1-pixel region. Its value is a ratio of the

- 53 -

features that are calculated for the cluster and the ring around the cluster. Feature 51 is the ratio of Feature 79 and Feature 23.

5 Feature 52 is the directional disparity of the 5x1 dark edge. Feature 52 is like Feature 51 except that its ratio is not of the 5x1 magnitude. Rather, it is of the 5x1 directional disparities of the cluster and ring around the cluster.

10 Feature 55 is called 5x5 dark edge strength. This is the ratio of Feature 27 and Feature 83, both of which were just described.

Feature 58 is called 3x3 blur residue mean. This is the ratio of Feature 30 and Feature 86, also just described.

15 Feature 64 is 7x7 blur residue skewness. This is the ratio of the 7x7 blur residue skewness of the cluster to that of the ring around the cluster. This ratio indicates how texture variations compare from the cluster to the ring around the cluster. It identifies changes in texture from the cluster to the ring around the cluster. It also provides a measure of how similar the cluster is to its background when its spatial frequencies have been eliminated by a 7x7 filter.

25 Whole image measurements are the second most diverse category of features. (A whole image is equivalent to one 20x FOV.) These features represent a third way that contextual information can be measured for a cluster. Several of these features are provided by the single cell process; others are calculated during the segmentation portion of the group process.

30 Feature 146 is the single cell process's high_count variable measurement, which measures dark edge gradient content of the whole original image.

- 54 -

This is a measure of how much cellular material may be in the image.

Feature 147 is the single cell process's high_mean variable measurement. This is the average value of all pixels in an image that have values between 199 and 250. This feature provides some information about an image's background.

Feature 151 is the "edge" variable determined during group segmentation. It is a combination of a 3x3 dilation residue and a 3x3 blur residue which is averaged over the whole image. This feature measures high spatial frequency and dark edge content of an image. It gives an indication of how much detail the image contains.

Feature 152 is the magnitude of the 2x1 dark edge in an image. This feature is calculated in the same way as Feature 21 except that it is calculated over the whole image.

Feature 153 is directional disparity of the 2x1 dark edge. This feature is calculated the same way as Feature 14 except that 1-pixel dark spots are searched for over the whole image.

Feature 154 is the magnitude of the 5x1 dark edge in an FOV. This feature is calculated the same way as Feature 152 except that dark 3-pixel regions are searched for instead of single dark pixels.

Feature 157 is 9x1 dark edge directional disparity. This is calculated the same way as Feature 153 except that 7-pixel dark regions are searched for rather than dark regions of single dark pixels.

Feature 158 is 5x5 dark edge strength. This feature is the same as Feature 27 except that it is calculated over the whole image.

Feature 160 is 17x17 dark edge strength. This feature is the same as Feature 29 except that the

- 55 -

feature is calculated over the whole image instead of just the cluster.

5 Feature 161 is 3x3 blur residue mean. This is calculated the same way as Feature 30 except that the feature is calculated over the entire image rather than just over the cluster.

10 Feature 162 is 3x3 blur residue standard deviation. This feature is measured the same way as Feature 31 except that it is calculated over the whole image.

Feature 165 is 7x7 blur residue mean. This feature is the same as Feature 161 except that lower spatial frequencies are measured over the whole image.

15 Feature 166 is 7x7 blur residue standard deviation. This is the same as Feature 91 except that it is calculated over the whole image.

20 Feature 169 is 15x15 blur residue mean. This feature is calculated the same way as Feature 165 except that a 15x15 binomial filter is used rather than a 7x7 binomial filter.

25 Feature 173 is standard deviation for the whole image. This is the standard deviation of the pixel values for every pixel in an image. This feature provides a measure of the amount of variation in pixel intensity there is across the entire 20x FOV.

30 Feature 174 is whole image skewness. This feature is the skewness of the values of every pixel in an image. It is a measure of how much pixel values are biased to one side or the other of mean pixel intensity. Skewness provides a measure of image content.

35 Feature 149 is the single cell process's low_threshold value. This value is calculated during single cell segmentation. It is the result of an adaptive threshold calculation for a certain range of

- 56 -

pixel intensities in an image. It gives a measure for how much dark matter there is in an image. If the threshold is low, there is a fair amount of dark matter in the image. If the threshold is high, there are probably few high density objects in the image.

The fifth and final step of image processing is called object classification. Each detected object is classified in an FOV according to values which were obtained during the feature extraction process.

Classification rules were defined during machine training.

Object classification is made by sets of statistical decisions through several levels of classification. Each classifier compares a set of features, unique to that classifier, measured from an object to decision boundaries which were determined during classifier training. Every segmented object is passed through filters and classifiers until the process reaches a classification determination. When a decision is made about an object, for example by one of the early box filters, results for that object (which are either artifact group or potential abnormal group) are reported. On the other hand, as long as a potential cell group cannot be identified, it continues down through filters and classifiers until a decision is reached regarding its classification. It may be useful to refer back to the five major steps of the endocervical group process's process described hereinabove.

Feature extraction (Step 4) and object classification (Step 5) are described separately, both occur simultaneously, over and over in sequence, until a decision is reached for a particular object. A segmented object from Step 3 enters the process at Step 4's first filter, which is the cytoplasm box

- 57 -

filter. Two features for that object are measured and extracted according to that filter's feature set (Features 2 and 106 in this case). If a determination can be reached by the first filter, processing for that object ceases, its results are reported, and feature extraction and object classification begin for the next segmented object in the current FOV. If a determination cannot yet be reached for an object, it is passed to the next processing level. Features are extracted at each classification level until the object is classified as an artifact, a probable endocervical cell group, a probable metaplastic/parabasal group, an uncertain endocervical/ metaplastic/parabasal group, or a potential abnormal cell group.

For each segmented object, feature extraction and object classification occur over and over until a determination is made. As soon as an object is classified, processing ceases at that level and the next object enters at level 1 for feature extraction and possible classification. The process's 14 layers of processing are:

1. cytoplasm box filter,
2. poly box filter,
- 25 3. miscellaneous box filter,
4. out of focus box filter,
5. graphite box filter,
6. bubble edge box filter,
7. 1st abnormal box,
- 30 8. 2nd abnormal box,
9. 3rd abnormal box,
10. stage 1 classifier,
11. stage 2 classifier,
12. stage 3 classifier,
- 35 13. post stage 3 box filter, and

- 58 -

14. object ID classifier.

The process tracks and accumulates results in each FOV so that classification of each object in an FOV is known. When the process has made a
5 determination for all objects in an FOV, results are reported to and accumulated by the system software and the 20x FOV integration process. Only then can image processing and classification begin for the next FOV.

After FOV segmentation is finished, objects are
10 subjected, one at a time, to a series of box filters. These simple classifiers perform threshold calculations using a pair of features.

The cytoplasm box filter is used to eliminate as many objects as possible that are cytoplasm
15 artifacts. This filter relies on two features which it calculates in the cluster area. One is a measure which compares texture of nuclei to that of the surrounding cytoplasm. The other feature is a measure of the variation in nuclear shapes. These two
20 features help identify objects where cytoplasm is incorrectly identified as nuclei in a significant number of cluster nuclei.

The purpose of the poly box filter is to classify and eliminate from further processing segmented
25 objects which are polymorphonucleocytes, commonly called polys. This filter uses two features which it calculates in the cluster area. The first feature measures the amount of small dark material in the cluster area. The second feature measures contrast,
30 which can be thought of as texture, in the cluster area. Polys have small dark nuclei which tend to be very dark when compared to their surroundings.

The filter removes many additional objects which clearly are not of interest. It uses two features
35 which it calculates in the cluster area. The first

- 59 -

feature calculates average nuclear perimeter, and the second measures high spatial frequency in the cluster.

The out of focus box filter is designed to remove objects that were poorly focused during image acquisition. Since identification of out of focus objects is unreliable, the process should not attempt to classify them. The first feature is measured in the cluster. It looks for high spatial frequencies. High spatial frequency is not typically observed in out of focus objects.

Glass laboratory slides of Pap smears commonly contain artifacts which prove to be graphite particles left by pencils. Since small particles of graphite can mimic nuclear clusters, it is useful to eliminate these artifacts early in processing. Because graphite particles are very dark, this filters two features measure contrast in the cluster area.

The endocervical group process may segment the edge of a bubble, mistaking it for an object. Segmented bubble edges appear as long, thin, dark objects. The first feature used by this filter is measured in the cluster and ring, then normalized. It looks for texture that is preferential in one direction, which is indicative of long thin objects. The second feature measures average minimum nuclear separation, which is generally quite small in segmented bubble edges.

For all objects that remain after the initial box filters, a check is made to determine if the objects are potentially abnormal, i.e. adenocarcinoma or high grade cell groups. If the first abnormal box identifies an object as potentially abnormal, that object is passed to the second abnormal box filter. Otherwise the object is passed to the Stage 1 classifier. The first and second abnormal box filters

- 60 -

make no final classification determinations. If an object is called potentially abnormal at the third abnormal box , processing ceases for that object and it is called potentially abnormal. Otherwise the object goes to the Stage 1 classifier.

Tables 4.1, 4.2 and 4.3 illustrate feature allocation for the three abnormal box filters. Table 4.1 illustrates how the first abnormal box filter's features are used. One density feature is measured in the cluster, and one density feature is measured over the whole image. One size feature is measured in the cluster area, and so on.

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	1	-	-	1
size	1	-	-	-
shape	-	-	-	-
texture	2	2	-	1
orientation	1	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	1

Table 4.1

- 61 -

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	-	-	-	1
size	3	-	-	-
shape	-	-	-	-
texture	3	2	-	1
orientation	-	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	1

10 Table 4.2

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	-	-	-	-
size	1	-	-	-
shape	-	-	-	-
texture	3	3	-	5
orientation	-	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	-

20 Table 4.3

Stage 1 is a Fisher's linear binary decision tree. Objects passed to the Stage 1 classifier should be normal or artifact groups. The Stage 1 classifier either identifies an object passed to it as an artifact, or it passes the object to the Stage 2 classifier for further refinement. Allocation of these features is illustrated in Table 4.4. See Table A.1 for more information about the features in these and the other feature sets.

- 62 -

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	1	-	1	1
size	1	-	-	-
shape	-	-	-	-
texture	4	-	2	-
orientation	1	-	-	-
N/C texture	-	-	-	-
N/C size	1	-	-	-
content	-	-	-	-

10 Table 4.4

Stage 2 is a Fishers linear binary decision tree. Objects passed to Stage 2 are either classified as artifacts, or they are passed to the Stage 3 classifier for further refinement. Table 4.5 illustrates how Stage 2's features are allocated.

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	1	-	-	1
size	1	-	-	-
shape	-	-	-	-
texture	1	-	-	7
orientation	3	1	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	-

25 Table 4.5

Stage 3 is a Fisher's linear binary decision tree. Objects passed to the Stage 3 classifier are either identified as artifacts, or they are passed to the post-Stage 3 box filter. Table 4.6

- 63 -

illustrates how Stage 3's features are allocated.

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	-	-	1	-
size	1	-	-	-
shape	-	-	-	-
texture	6	-	1	-
orientation	1	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	-

Table 4.6

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	-	-	-	-
size	-	-	-	-
shape	2	-	-	-
texture	-	1	-	2
orientation	-	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	-

Table 4.7

The object ID classifier is a Fisher's linear binary decision tree. Most objects passed to the object ID classifier are either endocervical cell groups or metaplastic/parabasal cell groups. This classifier identifies objects as probable endocervical, probable metaplastic/parabasal, or uncertain endocervical/metaplastic/parabasal.

- 64 -

	Cluster	ring around cluster	ring and cluster combined	whole image
CHARACTERISTICS				
density	2	-	-	-
size	-	-	-	-
shape	-	-	-	-
texture	4	1	-	-
orientation	1	-	-	-
N/C texture	-	-	-	-
N/C size	-	-	-	-
content	-	-	-	-

10 Table 4.8

When the endocervical group process has finished processing a slide, it has made a decision about each FOV and about every segmented object in each FOV. A slide's FOVs have been classified as:

- 15
1. could not contain cell groups,
 2. does not contain cell groups, or
 3. contains potential cell groups.

Likewise, each object in every FOV has been classified as:

- 20
1. an artifact group,
 2. a probable endocervical cell group,
 3. a potentially abnormal cell group,
 4. a probable metaplastic/parabasal cell group, or
 - 25 5. an uncertain endocervical/metaplastic/parabasal cell group.

As processing for each FOV is finished, results of all its objects are passed to the 20x FOV integration process. The integration process accumulates results, object by object. When results

30

- 65 -

are finally accumulated for all objects in every FOV, processing for that slide is complete. When each FOV is finished, its results are passed to the thick group process, the 20x FOV integration process, and to system software. System software accumulates a slide's results and conveys them to the scoring processes.

Process performance has been tracked and validated throughout all stages of classification training. A cross validation method was adapted for performance tracking at each stage, in which training data is randomly divided into ten equal sets. A classifier is then trained by nine of the ten sets and tested on the remaining set. Sets are rotated and the process is repeated until every combination of nine sets has been used for testing.

	density	7, 111, 116, 119		48	146, 147
	size	5, 103, 104, 125			
	shape	6, 99, 106			
	texture	8, 13, 14, 15, 18, 20, 21, 25, 27, 28, 29, 30, 31, 121, 136	77, 81, 83, 84, 86, 88, 91	51, 52, 55, 58, 64	151, 152, 153, 154, 157, 158, 160, 161, 162, 165, 166, 169, 173, 174
5	orien- tation	9, 42, 101	72		
	nuclear/ cytoplasm texture	2			
10	nuclear/ cytoplasm size	3			
	content				149

Table A.1 Features, by Type (vertical axis), as Measured in Different Areas (horizontal axis) of the Image

The invention has been described herein in considerable detail in order to comply with the Patent Statutes and to provide those skilled in the art with the information needed to apply the novel principles and to construct and use such specialized components as are required. However, it is to be understood that the invention can be carried out by specifically different equipment and devices, and

- 67 -

that various modifications, both as to the equipment details and operating procedures, can be accomplished without departing from the scope of the invention itself.

5

What is claimed is:

- 68 -

CLAIMS

1. A biological specimen classifier for classifying biological specimens on a slide with a slide coverslip, the apparatus comprising:
 - (a) a means for locating boundaries (12) of the slide coverslip;
 - (b) a means for image capture (12) having an image output representing at least one image of a biological specimen;
 - (c) a plurality of image classifiers (18, 20) wherein each of the plurality of image classifiers (18, 20) operates on the at least one image at a predetermined image magnification;
 - (d) a means for image data accumulation and slide scoring (22) connected to an output of at least one of the plurality of classifiers (18, 20) having a slide score and information output; and
 - (e) means for slide score reporting and information archival storage (26) connected to the slide score and information output.
2. A method of processing biological specimens on a slide with a slide coverslip comprising the steps of:
 - (a) locating boundaries of the slide coverslip (28);
 - (b) locating a biological specimen inside the boundaries of the slide coverslip (30);
 - (c) performing a first scan of the slide at a predetermined magnification to extract slide feature information (32);

- 69 -

- (d) accumulating information at the predetermined magnification (34);
 - (e) performing a second scan of the slide at a higher magnification including at least one field of view (38); and
 - (f) accumulating information and at least one slide score for archival storage (40).
3. The method of claim 2 wherein the first scan (32) comprises:
- (a) steps of image capture (44); and
 - (b) an information extraction step (46, 48).
4. The method of claim 2 wherein the second scan (38) comprises:
- (a) steps of image capture (56); and
 - (b) an information extraction step (58, 60).
5. A method of processing information at a low magnification scan comprising the steps of:
- (a) acquiring an image of a biological specimen to obtain a first image (44);
 - (b) performing at least one information extraction step on the acquired image (46, 48);
 - (c) compiling image information extraction results (50);
 - (d) repeating steps (a) and (b) until a predetermined sequence of images have been acquired; and
 - (e) selecting at least one image for further processing (52).
6. The method of claim 5 wherein the step of (e), further processing (52) comprises operating at

- 70 -

least one classifier (20).

7. A method of classifying objects in a biological specimen comprising the steps of:
- (a) obtaining at least one image (5) of the biological specimen;
 - (b) classifying (8) the at least one image according to at least one feature extraction classifier (7) to provide at least one classified object feature output;
 - (c) classifying a slide on a basis of the at least one classified object feature output to provide a slide classification output (9).
8. A method of classifying a biological specimen slide comprising the steps of:
- (a) performing a check on illumination aberrations caused by inappropriate obstructions in an illumination path (58, 60);
 - (b) performing a free-lying cell and isolated nucleus classification to identify objects as an intermediate cell, a normal cell, an artifact or a potentially abnormal cell and providing object classification, stain evaluation, object feature results and error information output (58, 60);
 - (c) processing an image to detect thick groups of cells (58, 60); and
 - (d) processing the image to process groups of cells (58, 60).
9. A biological specimen slide scoring apparatus

- 71 -

comprising

- 5 (a) a user interface means for receiving commands (542), storing and archiving data having at least one field of view results input, a slide score input, an endocervical score input, a cell count input, and a suitability input, wherein the user interface means has a report output;
- 10 (b) a means for image processing (46, 48) wherein the image processing means (46, 48) has an image output and an accumulated field of view results output (50); and
- 15 (c) a processing means (542) for processing images of at least one biological specimen slide having an imaging input connected to the image output (72), an accumulated field of view results input connected to the accumulated field of view results output (76) and the processing means (542)
- 20 has a field of view results output connected to the field of view results input (74), a slide score output connected to the slide score input, an endocervical score output connected to the endocervical score input, cell count output connected to the cell count input, and a suitability output connected to the suitability input (78).
- 25
- 30 10. The apparatus of claim 9 wherein the processing means includes a work station interface (542).
11. The apparatus of claim 9 wherein the processing means includes a computer database (70).

- 72 -

12. The apparatus of claim 9 further comprising a means for training (Fig. 9) at least one classifier.
13. An object classifier comprising a:
- 5 (a) means for identifying at least one object (102) in a biological specimen slide;
- (b) means for discriminating features (104) in the at least one identified object (102); and
- 10 (c) means for analyzing the features (106) to determine a class to which the at least one identified object (102) belongs.
14. A classifier training apparatus comprising a:
- 15 (a) means for obtaining an image (96) of a biological specimen slide;
- (b) means for identifying cells (98) on the biological specimen slide;
- (c) means for obtaining objects (102) of cells on a biological specimen slide;
- 20 (d) means for extracting features (104) of the cells and objects;
- (e) means for storing (114) a predetermined list of features;
- (f) means for classifying (106) the features based on a selection of features; and
- 25 (g) means for evaluating (110) an output of a classifier.
15. A trainable object classifier comprising a:
- 30 (a) means for obtaining an image (100) of a biological specimen slide (96) that have been predetermined as representing a training set;

- 73 -

- (b) a means for identifying objects (102) in the image;
 - (c) means for classifying (106) a set of objects based on a predetermined classification object feature set (114).
- 5
16. A method of building a classifier for a biological specimen analysis system comprising the steps of:
- (a) inputting a semi-automatic cell library into an initial classifier (146);
 - (b) running a second classifier (148) on an output of the initial classifier (146) in response to a user defined set of features (144) wherein the second classifier (148) has a classifier result output; and
 - (c) repeating step (a) and (b) until classification results meet a predefined objective.
- 10
- 15
17. A method of developing a biological specimen classifier comprising the steps of:
- (a) obtaining images (162) of a plurality of biological specimen slides;
 - (b) classifying (164) the plurality of biological specimen slides according to a first predetermined classification;
 - (c) evaluating (166) the output of the classifier according to predetermined evaluation criteria (186, 188, 190);
 - (d) determining if the output of the evaluation determines that the classification was satisfactory (168) according to predetermined criteria (186, 188, 190);
- 20
- 25
- 30

- 74 -

- (e) reclassifying the plurality of slides on a second classifier (172); and
- (f) repeating steps (b) - (e) until the output of the second classifier (172) meet
5 predetermined criteria (186, 188, 190).

18. A biological specimen classifier for classifying biological specimens on a slide with a slide coverslip, the apparatus comprising:

- 10 (a) classifier means (196), connected to receive at least one image of a biological specimen (194), for classifying the biological specimen having at least one classification result output (200) and at
15 least one test output (202); and
- (b) means, connected to the classification result output (200) and the at least one test output (202), for determining if the classification result is valid (206).

20 19. The apparatus of claim 18 wherein the classifier means (196) further comprises means for performing a system integrity test.

20. The apparatus of claim 18 wherein the classifier means (196) further comprises means
25 for testing for an acceptable coverslip size and location.

21. The apparatus of claim 18 wherein the classifier means (196) further comprises means
30 for testing for an acceptable thickness of a coverslip and mounting medium.


- 75 -

22. The apparatus of claim 18 wherein the classifier means (196) further comprises a means for testing for an acceptable quantity of specimen.
- 5 23. The apparatus of claim 18 wherein the classifier means (196) further comprises a means for testing for acceptable focus quality.
24. The apparatus of claim 18 wherein the classifier means (196) further comprises a
10 means testing for presence of bubbles in a mounting medium.
25. The apparatus of claim 18 wherein the classifier means (196) further comprises a
15 means for testing for proper staining and preparation of a specimen.
26. The apparatus of claim 18 wherein the classifier means (196) further comprises means for test for specimen abnormality.
27. A method of validating the classification
20 results of a biological specimen classifier for classifying biological specimens on a slide with a slide coverslip, the apparatus comprising:
- 25 (a) obtaining an image (100) of the biological specimen;
- (b) performing a system integrity test (196);
- (c) testing for an acceptable coverslip size and location (196);
- 30 (d) testing for an acceptable thickness of the coverslip and mounting medium (196);


- 76 -

- (e) testing for an acceptable quantity of specimen (196);
- (f) testing for acceptable focus quality (196);
- 5 (g) testing for the presence of bubbles in the mounting medium (196);
- (h) testing for proper staining and preparation of the specimen (196); and
- 10 (i) determining the validity of the classification results (206).


1/16



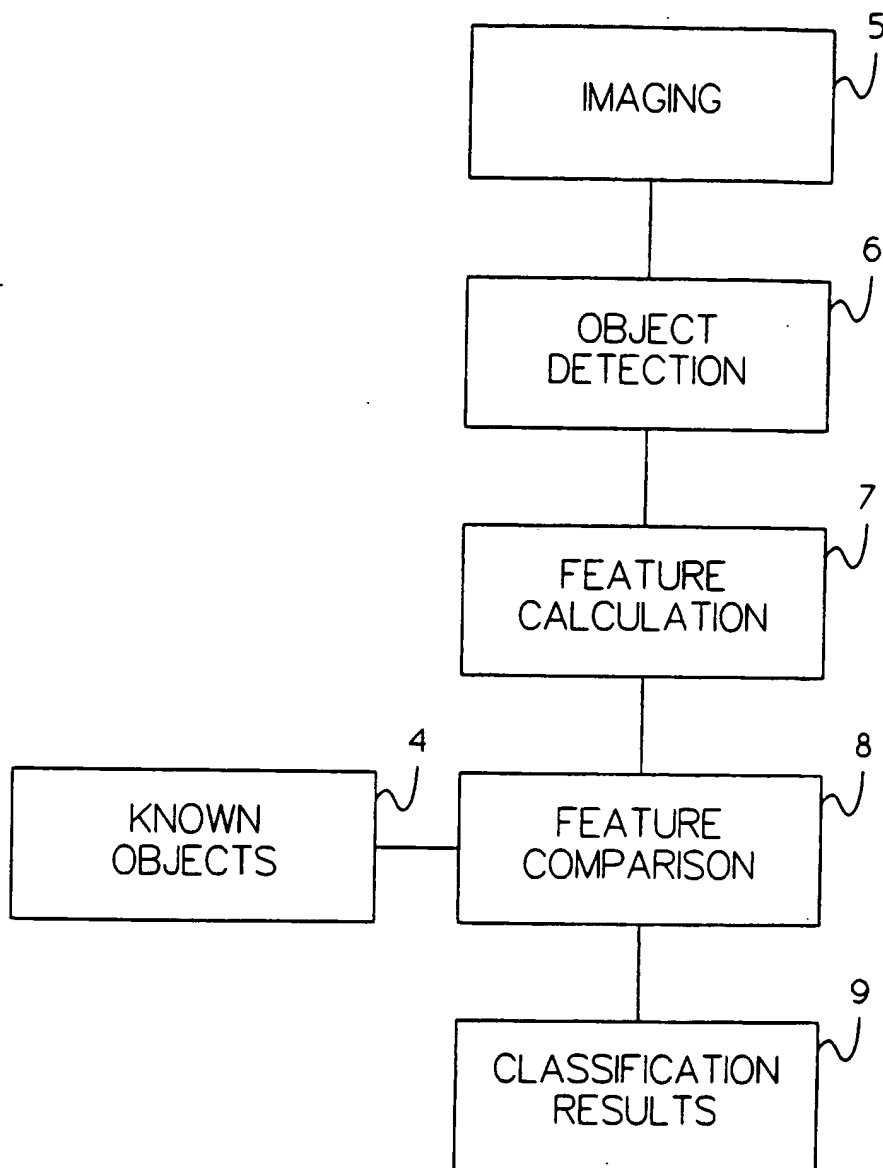
2/16



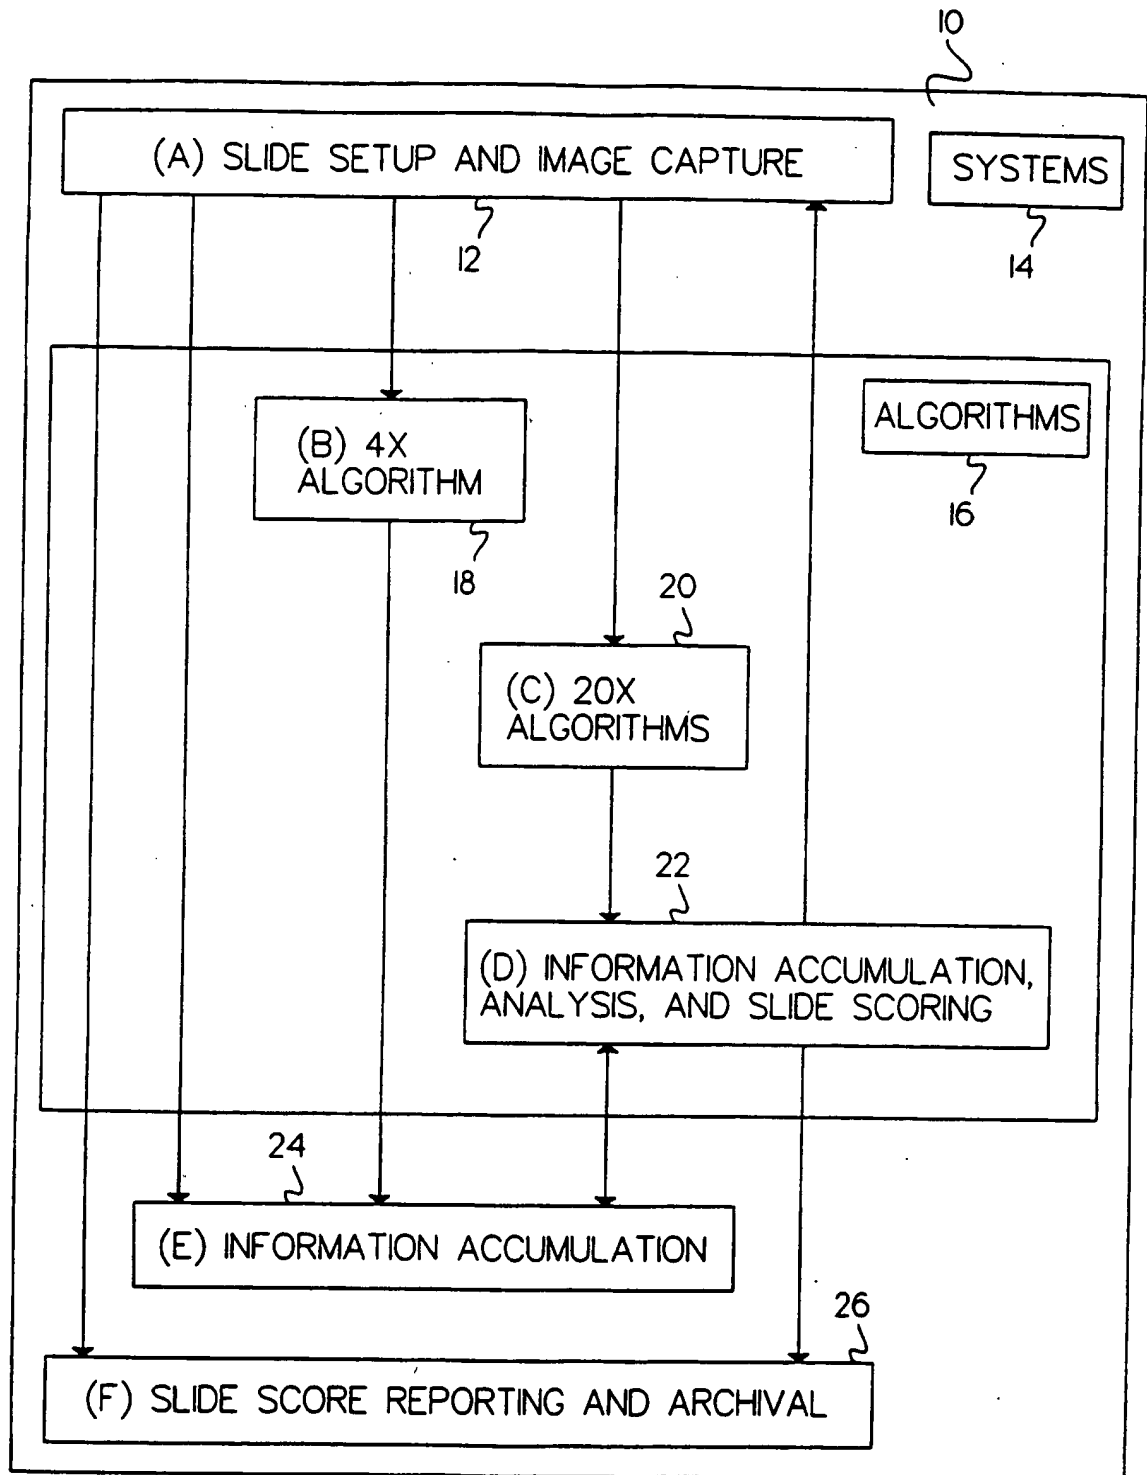
3/16



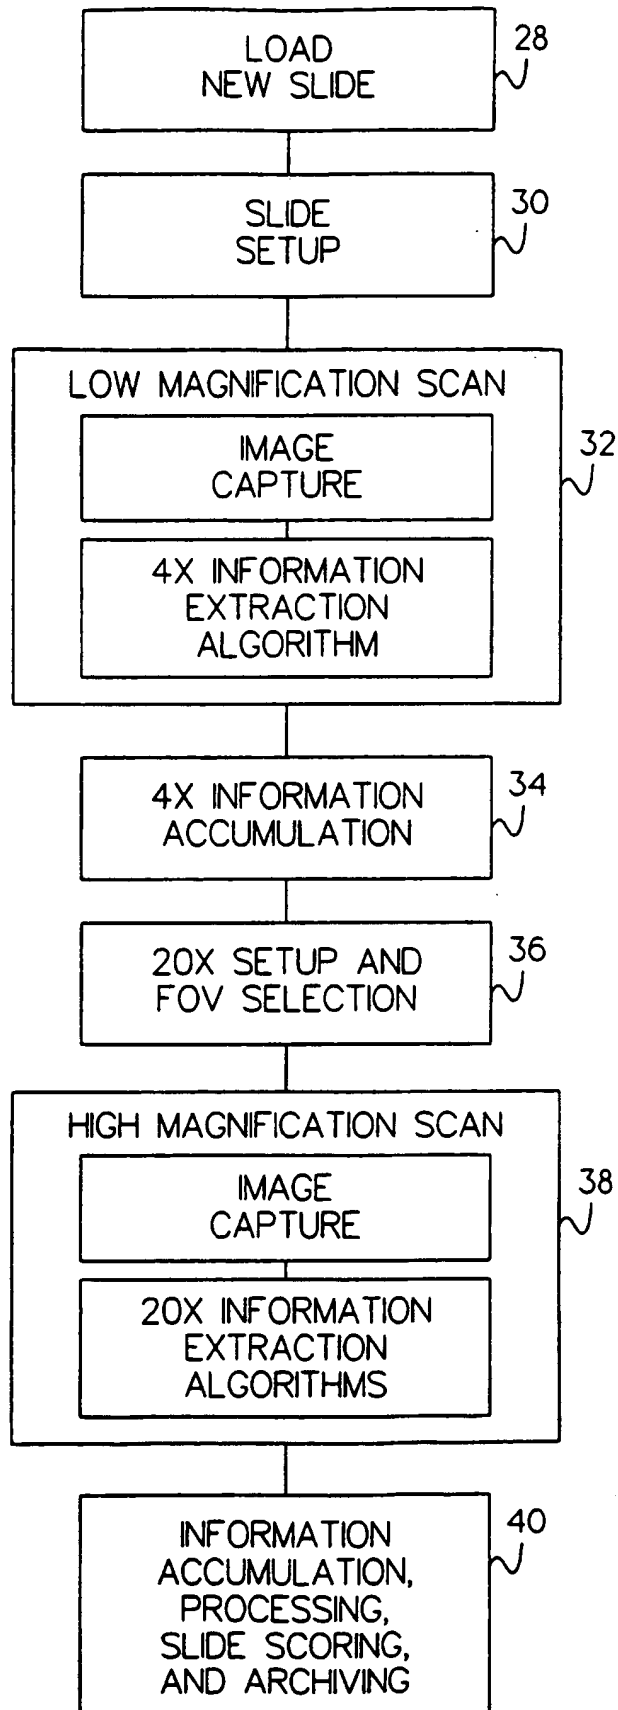
4/16

Fig. 2A


5 / 16

Fig.- 2B

6/16

Fig- 3

7/16



8/16


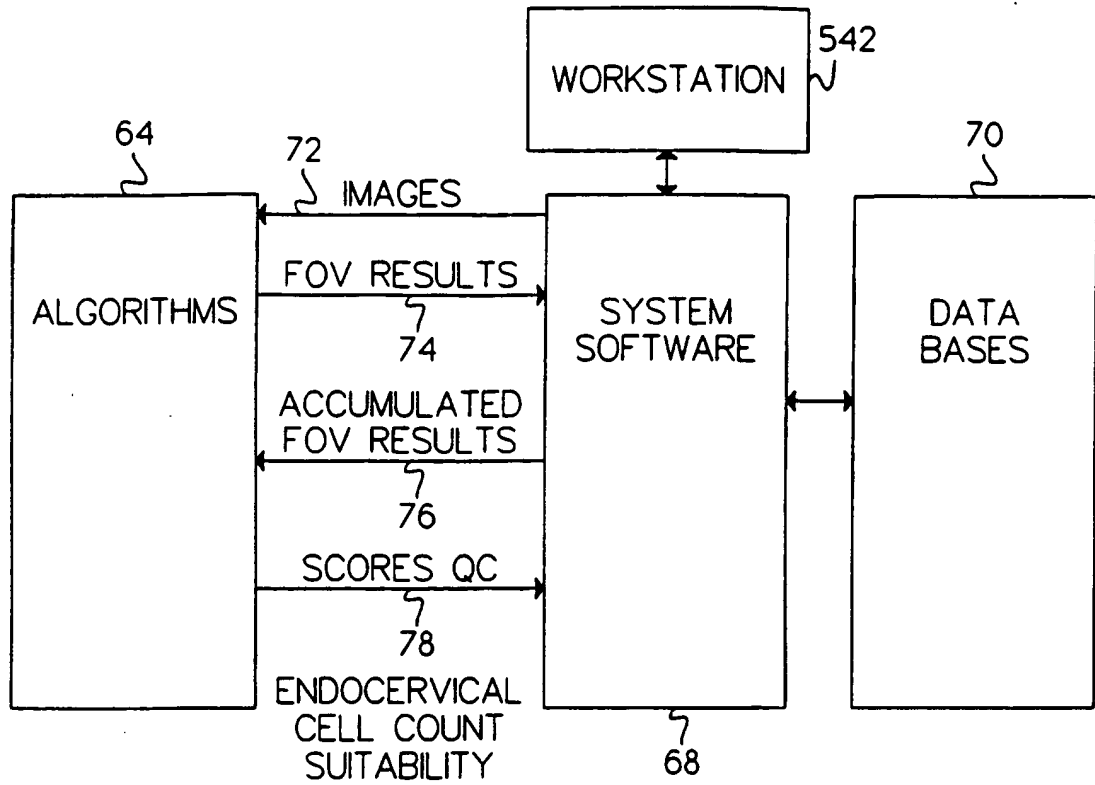



Fig-5

9/16

Fig- 6Fig- 7

10 / 16






Fig- 8


11/16

Fig- 9


12/16

Fig-10

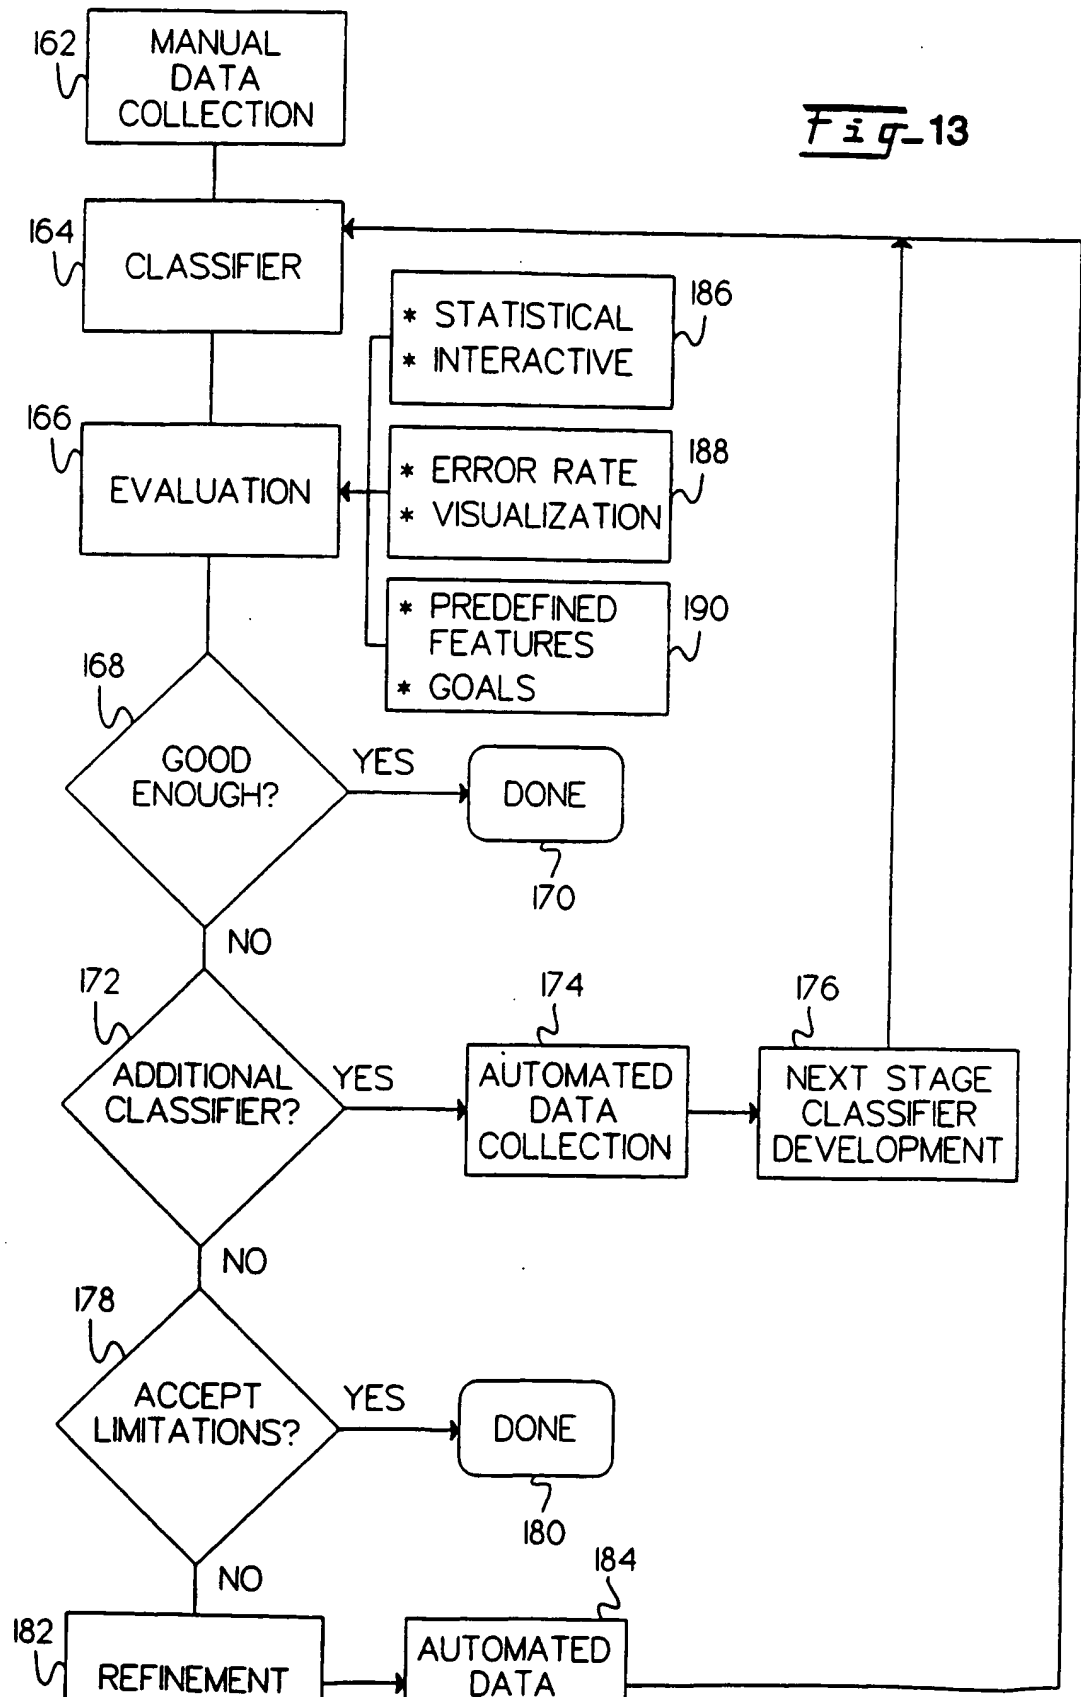
13/16

Fig-11


14/16

Fig-12

15 / 16

Fig-13

16 / 16

Fig-14

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/10165

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06K 9/00

US CL : Please See Extra Sheet.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 382/128, 129, 130, 132, 133, 134; 356/ 39, 40; 377/10, 11, 12, 13

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, Search terms: classifier, boundaries, magnification

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US, A, 5,235,522 (BACUS) 10 AUGUST 1993, col. 5, line 43 to col. 16, line 7; Figs. 1-14	1 --- 2-27
X -- Y	US, A, 4,175,860 (BACUS) 27 NOVEMBER 1979, col. 4, line 31 to col. 12 line 7, figs. 1, 2, 4, 5	2-4 ----- 1, 5-27
Y	US, A, 5,121,320 (AOKI ET AL.) 09 JUNE 1992, col. 4, line 42 to col. 13, line 37; fig. 2	14-17
Y	US, A, 5,202,931 (BACUS) 13 APRIL 1993, whole document	1-27
Y	US, A, 4,741,043 (BACUS) 26 APRIL 1988, whole document	1-27

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be part of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"A"	document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

05 DECEMBER 1995

Date of mailing of the international search report

22 JAN 1996

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3220

Authorized officer
[Signature]
BIPIN SHALWALA

Telephone No. (703) 305-9576

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/10165

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 5,281,517 (BACUS ET AL.) 25 JANUARY 1994, whole document	1-27
Y	US, A, 5,018,209 (BACUS) 21 MAY 1991, whole document	1-27
Y	US, A, 5,257,182 (LUCK ET AL.) 26 OCTOBER 1993 whole document, specifically fig. 2	2-4
Y,P	US, A, 5,361,140 (HAYENGA ET AL.) 01 NOVEMBER 1994, Fig. 2, col. 4, lines 20-58	2-4
Y	US, A, 5,287,272 (RUTENBERG ET AL.) 15 FEBRUARY 1994 whole document	1-27

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/10165

A. CLASSIFICATION OF SUBJECT MATTER:

US CL :

382/128, 129, 130, 132, 133, 134; 356/ 39, 40; 377/10, 11, 12, 13